Screening Tests for Lasso Problems

Zhen James Xiang, Yun Wang, and Peter J. Ramadge

Abstract—This paper is a survey of dictionary screening for the lasso problem. The lasso problem seeks a sparse linear combination of the columns of a dictionary to best match a given target vector. This sparse representation has proven useful in a variety of subsequent processing and decision tasks. For a given target vector, dictionary screening quickly identifies a subset of dictionary columns that will receive zero weight in a solution of the corresponding lasso problem. These columns can be removed from the dictionary, prior to solving the lasso problem, without impacting the optimality of the solution obtained. This has two potential advantages: it reduces the size of the dictionary, allowing the lasso problem to be solved with less resources, and it may speed up obtaining a solution. Using a geometrically intuitive framework, we provide basic insights for understanding useful lasso screening tests and their limitations. We also provide illustrative numerical studies on several datasets.

Index Terms—sparse representation, feature selection, lasso, dual lasso, dictionary screening.

arXiv:1405.4897v1 [cs.LG] 19 May 2014

1 Introduction

Recently, the sparse representation of data with respect to a dictionary of features has contributed to successful new methods in machine learning, pattern analysis, and signal/image processing. At the heart of many sparse representation methods is the least squares problem with ℓ_1 regularization, often called the lasso problem [1]:

$$\min_{\mathbf{w}\in\mathbb{R}^p} \quad 1/2\|\mathbf{y}-B\mathbf{w}\|_2^2+\lambda\|\mathbf{w}\|_1, \tag{1}$$

where $\lambda>0$ is a regularization parameter. The matrix $B\in\mathbb{R}^{n\times p}$ is called the *dictionary* and its columns $\{\mathbf{b}_i\}_{i=1}^p$ are usually called *features*. Depending on the field, the terms *codewords*, *atoms*, *filters*, and *regressors* are also used. The lasso problem seeks a representation of the *target vector* $\mathbf{y}\in\mathbb{R}^n$ as a linear combination $\sum_{i=1}^p w_i\mathbf{b}_i$ of the features with many $w_i=0$ (sparse representation). Equation (1) also serves as the Lagrangian for the widely used constrained problems $\min_{\mathbf{w}\in\mathbb{R}^p}\ \|\mathbf{y}-B\mathbf{w}\|_2^2$ subject to $\|\mathbf{w}\|_1\le\sigma$, and $\min_{\mathbf{w}\in\mathbb{R}^p}\ \|\mathbf{w}\|_1$ subject to $\|\mathbf{y}-B\mathbf{w}\|_2^2\le\varepsilon$. Many solvers of these problems address the Lagrangian formulation (1) directly [2].

The above problems are studied extensively in the signal processing, computer vision, machine learning, and statistics literature. See, for example, the general introduction to sparse dictionary representation methods in [3] and [4]. Sparse representation has proven effective in applications ranging from image restoration [5], [6], to face recognition [7], [8], object recognition [9], speech classification [10], speech recognition [11], music genre classification [12], and topic detection in text documents [13]. In these applications, it is common to encounter a large dictionary (e.g., in face recognition), data with large data dimension (e.g., in topic detection), and in dictionary learning, a large number of dictionary iterations (e.g., in image restoration). These factors can make solving problem (1) a bottleneck in the computation.

Several approaches have been suggested for addressing this computational challenge. In the context of classification, Zhang et al. [14] propose abandoning sparsity and using a fast collaborative linear representation scheme based on ℓ_2 regularized least squares. This improves the speed of classification in face recognition applications. However, Xu and Ramadge [15] point out that in general collaborative representation is not always required, and the (nonlinear) Sparse Representation Classifier (SRC) [7] achieves superior classification accuracy. Another approach is to seek a sparse representation using a fast greedy method to approximate the solution of (1). There has been a considerable amount of work in this direction, see for example [3], [16], [17]. However, this approach seems best when seeking very sparse solutions and, in general, the solutions obtained can be challenging to analyze.

Recently an approach known as (dictionary) screening has been proposed. For a given target vector $\mathbf{y}$ and regularization parameter λ, screening quickly identifies a subset of features that is guaranteed to have zero weight in a solution $\hat{\mathbf{w}}$ of (1). These features can be removed (or "rejected") from the dictionary to form a smaller, more readily solved lasso problem. By padding its solution appropriately with zeros, one can obtain a solution of the original problem. This approach is the focus of the paper.

Screening has two potential benefits. First, it can be run in an on-line mode with very few features loaded into memory at a given time. By this means, screening can significantly reduce the size of the dictionary that needs to be loaded into memory in order to solve

• *Authors are with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544, USA*
E-mail: {zxiang, yunwang, ramadge}@princeton.edu

the lasso problem. Second, by quickly reducing the number of features we can often solve problems faster. Even small gains can become very significant when many lasso problems must be solved. Moreover, since screening is transparent to the lasso solver, it can be used in conjunction with many existing solvers.

The idea of screening can be traced back to various feature selection heuristics in which selected features $\{\mathbf{b}_i\}$ are used to fit a response vector $\mathbf{y}$. This is usually done by selecting features based on an empirical measure of relevance to $\mathbf{y}$, such as the correlation of $\mathbf{y}$ and $\mathbf{b}_i$. This is used, for example, in univariate voxel selection based on t-statistics in the fMRI literature [18]. Fan and Lv [19] give an excellent review of recent results on correlation based feature selection and formalize the approach in a probabilistic setting as a correlation based algorithm called Sure Independence Screening (SIS). In a similar spirit, Tibshirani et al. [20] reports Strong Rules for screening the lasso, the elastic net and logistic regression. These rules are also based on thresholding correlations. With small probability, SIS and the Strong Rules can yield "false" rejections.

A second approach to screening seeks to remove dictionary columns while avoiding any false rejections. In spirit, this harks back to the problem of removing "non-binding" constraints in linear programs [21]. For the lasso problem, the first line of recent work in this direction is due to El Ghaoui et al. [22], where such screening tests are called "SAFE" tests. In addition to the lasso, this work examined screening for a variety of related sparse regularization problems. Recent work ([23], [24], [25], [26], [27]) has focused mainly on the lasso problem and close variants.

The basic approach in the above papers is to bound the solution of the dual problem of (1) within a compact region $\mathcal{R}$ and find $\mu_{\mathcal{R}}(\mathbf{b}) = \max_{\boldsymbol{\theta}\in\mathcal{R}} \boldsymbol{\theta}^T\mathbf{b}$. For simple regions $\mathcal{R}$, $\mu_{\mathcal{R}}$ is readily computed and yields a screening test for removing a subset of unneeded features. This approach has resulted in tests based on spherical bounds [22], [23], the intersection of spheres and half spaces (domes) [22], [24], elliptical bounds [26] and novel approaches for selecting the parameters of these regions to best bound the dual solution of (1) [25], [27]. These screening tests can execute quickly, either serially or in parallel, and require very few features to be loaded into memory at once. If one seeks a strongly to moderately sparse solution, the tests can significantly reduce dictionary size and speed up the solution of lasso problems.

To keep our survey focused, we concentrate on screening for the lasso problem. However, the methods discussed apply to any problem that can be efficiently transformed into a lasso problem. For example, the elastic net [28] and full rank generalized lasso problems [29]. Moreover, the basic ideas and methods discussed are a good foundation for applying screening to other sparse regularization problems. We will situate our exposition within the context of prior work as the development proceeds.

The main features of our survey include:
(a) Our exposition uses a geometric framework which unifies many lasso screening tests and provides basic tools and geometric insights useful for developing new tests. In particular, we emphasize the separation of the structure or "architecture" of the test from the design problem of selecting its parameters.
(b) We examine whether more complex screening tests are worthwhile. For each $m \geq 0$, there is a family of tests based on the intersection of a spherical bound and m half spaces. As m increases these tests can reject more features but are also more time consuming to execute. To examine if more complex tests are worthwhile, we derive the region screening test for the intersection of a sphere and two half spaces, and use this to examine where current region screening tests stand in the trade-off between rejection rate and computational efficiency.
(c) We show how composite tests can be formed from existing tests. In particular, we describe a composite test based on carefully selected dome regions that performs competitively in numerical studies. We also point out a fundamental limitation of this approach.
(d) We review sequential screening schemes that make headway on the problem of screening for small normalized values of λ. When used in an "on-line" mode with realistic values of the regularization parameter, these methods can successfully reduce the size of large dictionaries to a manageable size, and allow lasso problems to be solved faster than without screening.

1.1 Outline of the Paper

We begin in §2 with a review of basic tools, especially the dual of the lasso problem and its geometric interpretation. §3 introduces screening in greater detail and §4 introduces region tests. After these preparations, we discuss several important forms of region tests: sphere tests (§4.2), sphere plus hyperplane tests (§4.3) and sphere plus two hyperplane tests (§4.5). We show how spherical bounds can be iteratively refined using features (§4.4), and examine ways to combine basic tests. §5 gives a brief overview of sequential screening. We give a practical summary of screening algorithms in §6 and illustrate the results of screening via numerical studies in §7. We conclude in §8. Proofs of new or key results are given in the Appendices, organized by the section in which the result is discussed.

2 Preliminaries

We focus on the lasso (1), but it will be convenient to also consider the *nonnegative* lasso:

$$\begin{array}{ll} \min\limits_{\mathbf{w}\in\mathbb{R}^p} & {}^1\!/_2\|\mathbf{y} - B\mathbf{w}\|_2^2 + \lambda\|\mathbf{w}\|_1, \\ \text{s.t.} & \mathbf{w} \geq 0. \end{array} \tag{2}$$

The analysis and algorithms in the paper apply (with minor changes) to both problems.

Throughout the paper we assume that a fixed dictionary B is used to solve various instances of (1) or (2). We assume that all features are nonzero and say that the dictionary is *normalized* if all features have unit norm. Each instance is specified by a pair $(\mathbf{y}, \lambda)$ consisting of a target vector $\mathbf{y}$ and a value λ of the regularization parameter.

Multiplying the objective of (1) by α^2, with $\alpha > 0$, yields the equivalent problem:

$$\min_{\mathbf{w}\in\mathbb{R}^p} \ {}^1\!/_2\|\bar{\mathbf{y}} - \bar{B}\mathbf{w}\|_2^2 + \bar{\lambda}\|\mathbf{w}\|_1,$$

where $\bar{\mathbf{y}} = \alpha\mathbf{y}$, $\bar{B} = \alpha B$, and $\bar{\lambda} = \alpha^2\lambda$. Some lasso solvers require that $\|B\|_F \le 1$, and problem (1) must be scaled to ensure this holds. As a result, it is meaningless to talk about the value of λ employed when solving (1) without accounting for possible scaling. One way to do this is to define $\lambda_{\max} = \max_{j=1}^p |\mathbf{b}_j^T\mathbf{y}|$. Then the ratio $\lambda/\lambda_{\max}$ is invariant to scaling. The parameter $\lambda_{\max}$ is also useful for other purposes related to screening. Throughout the paper, we use the ratio $\lambda/\lambda_{\max}$ as an unambiguous measure of the amount of regularization used in solving (1) and (2).

Geometric insight on lasso problems, and on screening in particular, is enhanced by bringing in the Lagrangian dual of (1). The following parameterization of the dual problem is particularly convenient [23]:

$$\begin{aligned} \max_{\boldsymbol{\theta}\in\mathbb{R}^n} \quad & {}^1\!/_2\|\mathbf{y}\|_2^2 - {}^{\lambda^2}\!/_2\|\boldsymbol{\theta} - \mathbf{y}/\lambda\|_2^2 \\ \text{s.t.} \quad & |\boldsymbol{\theta}^T\mathbf{b}_i| \le 1 \quad \forall i = 1, 2, \ldots, p. \end{aligned} \tag{3}$$

Solutions $\hat{\mathbf{w}} \in \mathbb{R}^p$ of (1) and $\hat{\boldsymbol{\theta}} \in \mathbb{R}^n$ of (3) satisfy:

$$\mathbf{y} = B\hat{\mathbf{w}} + \lambda\hat{\boldsymbol{\theta}}, \quad \hat{\boldsymbol{\theta}}^T\mathbf{b}_i = \begin{cases} \operatorname{sign}\hat{w}_i, & \text{if } \hat{w}_i \neq 0; \\ \gamma \in [-1, 1]\,, & \text{if } \hat{w}_i = 0. \end{cases} \tag{4}$$

The corresponding dual problem of (2) is:

$$\begin{aligned} \max_{\boldsymbol{\theta}\in\mathbb{R}^n} \quad & {}^1\!/_2\|\mathbf{y}\|_2^2 - {}^{\lambda^2}\!/_2\|\boldsymbol{\theta} - \mathbf{y}/\lambda\|_2^2 \\ \text{s.t.} \quad & \boldsymbol{\theta}^T\mathbf{b}_i \le 1 \quad \forall i = 1, 2, \ldots, p, \end{aligned} \tag{5}$$

with the primal and dual solutions related via:

$$\mathbf{y} = B\hat{\mathbf{w}} + \lambda\hat{\boldsymbol{\theta}}, \ \hat{\boldsymbol{\theta}}^T\mathbf{b}_i = \begin{cases} 1, & \text{if } \hat{w}_i > 0; \\ \gamma \in (-\infty, 1]\,, & \text{if } \hat{w}_i = 0. \end{cases} \tag{6}$$

A derivation of (3) and (4) is given in the Appendix. It will be convenient to define a *feature pool* $\mathfrak{B}$. For the lasso, $\mathfrak{B} = \{\pm\mathbf{b}_i\}_{i=1}^p$ and for the nonnegative lasso, $\mathfrak{B} = \{\mathbf{b}_i\}_{i=1}^p$. This allows the constraints in (3) and (5) to be stated as $\forall \mathbf{b} \in \mathfrak{B} : \boldsymbol{\theta}^T\mathbf{b} \le 1$.

For $\mathbf{x} \in \mathbb{R}^n$, let $P(\mathbf{x}) = \{\mathbf{z} : \mathbf{x}^T\mathbf{z} = 1\}$ denote the hyperplane in $\mathbb{R}^n$ that has unit normal $\mathbf{x}/\|\mathbf{x}\|_2$ and contains the point $\mathbf{x}/\|\mathbf{x}\|_2^2$. Let $H(\mathbf{x}) = \{\mathbf{z} : \mathbf{x}^T\mathbf{z} \le 1\}$ denote the corresponding closed half space containing the origin. So a constraint of the form $\mathbf{b}^T\boldsymbol{\theta} \le 1$ requires that $\boldsymbol{\theta}$ lies in the closed half space $H(\mathbf{b})$. Hence the set of feasible points $\mathcal{F}$ of the dual problems is the nonempty, closed, convex polyhedron formed by the intersection of the finite set of closed half spaces

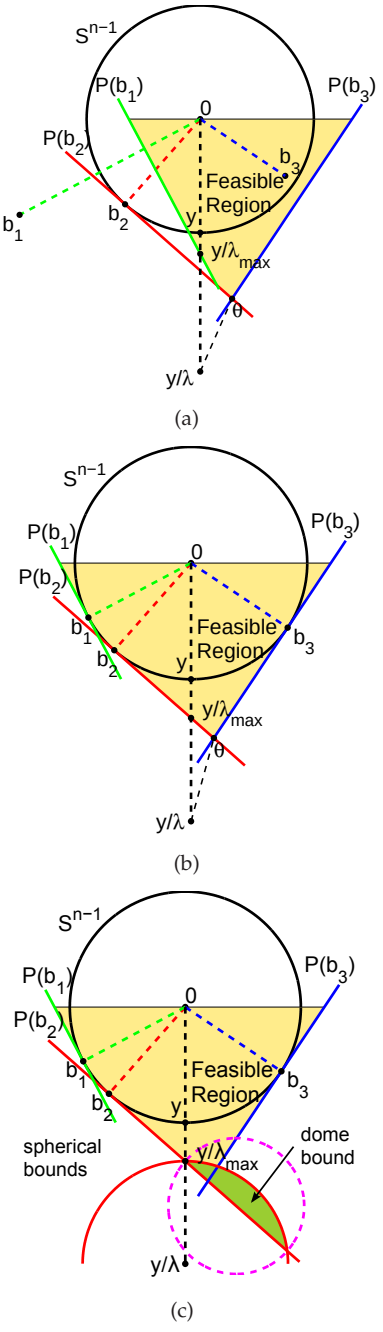

Fig. 1. The constraints and feasible set $\mathcal{F}$ of the dual problem for **(a):** general features, and **(b):** unit norm features. **(c):** Examples of two spheres and a dome region bounding $\hat{\boldsymbol{\theta}}$ when the features have unit norm. In all cases only the lower half of $\mathcal{F}$ is shown.

$H(\mathbf{b})$, $\mathbf{b} \in \mathfrak{B}$. This is illustrated in Fig. 1(a) and 1(b). In addition, for the lasso, $\boldsymbol{\theta} \in \mathcal{F}$ if and only if $-\boldsymbol{\theta} \in \mathcal{F}$. So $-\mathcal{F} = \mathcal{F}$. This follows from the same property of the feature pool: $-\mathfrak{B} = \mathfrak{B}$.

To maximize the objective function in (3) or (5) we seek the projection $\hat{\boldsymbol{\theta}}$ of $\mathbf{y}/\lambda$ onto the closed convex set $\mathcal{F}$. This is the *unique* point satisfying the following set of inequalities [30, §3.1]: for each $\boldsymbol{\theta} \in \mathcal{F}$,

$$(\mathbf{y}/\lambda - \hat{\boldsymbol{\theta}})^T(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}) \le 0. \tag{7}$$

In contrast, the lasso problem (1) may not have a unique solution [31], [32], [33].

The set of points $\{\hat{\boldsymbol{\theta}}(\lambda), \lambda > 0\}$ is called the *dual regularization path*. For λ sufficiently large, $\mathbf{y}/\lambda$ lies in $\mathcal{F}$ and $\hat{\boldsymbol{\theta}}(\lambda) = \mathbf{y}/\lambda$. To find the smallest λ for which this holds, let

$$\lambda_{\max} = \max_{\mathbf{b}\in\mathfrak{B}} \mathbf{y}^T\mathbf{b}, \tag{8}$$

$$\mathbf{b}_{\max} \in \arg\max_{\mathbf{b}\in\mathfrak{B}} \mathbf{y}^T\mathbf{b}. \tag{9}$$

Then for all $\mathbf{b} \in \mathfrak{B}$: $(\mathbf{y}/\lambda_{\max})^T\mathbf{b} \leq \mathbf{y}^T\mathbf{b}_{\max}/\lambda_{\max} = 1$. So $\mathbf{y}/\lambda_{\max}$ lies in the boundary of $\mathcal{F}$. As $\lambda \geq \lambda_{\max}$ decreases from a large value, $\hat{\boldsymbol{\theta}}(\lambda) = \mathbf{y}/\lambda$ moves in a straight line within $\mathcal{F}$ until $\lambda = \lambda_{\max}$, at which point $\hat{\boldsymbol{\theta}}(\lambda_{\max}) = \mathbf{y}/\lambda_{\max}$ first lies on the boundary of $\mathcal{F}$. As λ decreases below $\lambda_{\max}$, $\mathbf{y}/\lambda$ moves away from $\mathcal{F}$ and $\hat{\boldsymbol{\theta}}(\lambda)$ is the unique projection of $\mathbf{y}/\lambda$ onto the boundary of $\mathcal{F}$. Using (4), for $\lambda/\lambda_{\max} > 1$, $\hat{\mathbf{w}} = 0$, and conversely, if $\hat{\mathbf{w}} = 0$, then $\hat{\boldsymbol{\theta}} = \mathbf{y}/\lambda \in \mathcal{F}$. So for $\lambda/\lambda_{\max} \in (0, 1)$, $\mathbf{y}/\lambda \notin \mathcal{F}$, $\hat{\boldsymbol{\theta}}(\lambda)$ lines on the boundary of $\mathcal{F}$, and $\hat{\mathbf{w}}$ is nonzero.

Let $\mathcal{I} = [1, \ldots, p]$ denote the ordered set of feature indices and $S \subset \mathcal{I}$. Given $\mathbf{w} \in \mathbb{R}^p$, let $\mathbf{w}_{\downarrow S}$ denote the vector in $\mathbb{R}^{|S|}$ obtained by subsampling $\mathbf{w}$ at the indices in S. Conversely, for $\mathbf{z} \in \mathbb{R}^{|S|}$, let $\mathbf{z}^{\uparrow S}$ denote the vector in $\mathbb{R}^p$ obtained by upsampling $\mathbf{z}$: the entries of $\mathbf{z}^{\uparrow S}$ with indices in S take the corresponding values in $\mathbf{z}$ and all other entries are zero. Similarly, for a dictionary $B \in \mathbb{R}^{n\times p}$, let $B_{\downarrow S}$ denote the subdictionary obtained by sampling the columns of B at the indices in S. The following properties are clear: (a) $\mathbf{z} = (\mathbf{z}^{\uparrow S})_{\downarrow S}$; (b) if $w_i = 0$ for $i \notin S$, then $\mathbf{w} = (\mathbf{w}_{\downarrow S})^{\uparrow S}$; (c) $B_{\downarrow S}\mathbf{z} = B\mathbf{z}^{\uparrow S}$; and (d) if $w_i = 0$ for $i \notin S$, $B\mathbf{w} = B_{\downarrow S}\mathbf{w}_{\downarrow S}$.

By (4), if we know the primal solution $\hat{\mathbf{w}}$, then the dual solution is $\hat{\boldsymbol{\theta}} = (\mathbf{y} - B\hat{\mathbf{w}})/\lambda$. Conversely, if we know the dual solution $\hat{\boldsymbol{\theta}}$, then any point satisfying the following equations is a primal solution [33]:

$$\begin{aligned} B_{\downarrow A(\hat{\boldsymbol{\theta}})}\mathbf{w}_{\downarrow A(\hat{\boldsymbol{\theta}})} &= \mathbf{y} - \lambda\hat{\boldsymbol{\theta}} \\ \mathbf{w}_{\downarrow A(\hat{\boldsymbol{\theta}}),i}(\hat{\boldsymbol{\theta}}^T\mathbf{b}_i) &\geq 0,\ i \in A(\hat{\boldsymbol{\theta}}) \end{aligned} \tag{10}$$

where $A(\hat{\boldsymbol{\theta}}) = \{i\colon |\hat{\boldsymbol{\theta}}^T\mathbf{b}_i| = 1\}$.

3 Screening

We now explain the idea of screening in detail. Given an instance $(\mathbf{y}, \lambda)$ of (1), we select a partition $\mathcal{I} = S \cup \bar{S}$ of the features. We say that the features indexed by S are *selected* and those indexed by $\bar{S}$ are *rejected*. Then we form the reduced dictionary $B_{\downarrow S}$ of selected features and let $\hat{\mathbf{z}}$ denote a solution of the corresponding lasso problem using this dictionary. In general, the upsampled vector $\hat{\mathbf{z}}^{\uparrow S}$ is not a solution of the original lasso problem (1). Here is the key point: screening seeks a partition such that the upsampled vector $\hat{\mathbf{z}}^{\uparrow S}$ solves (1). In general, such a partition depends on the instance and hence must be computed "on-the-fly".

By virtue of being smaller, the reduced problem is more manageable. For example, it may fit into memory when the original problem does not, and finding its solution may require less time. Hence there are two evaluation metrics of interest: the size of S (or $\bar{S}$) as a fraction of $\mathcal{I}$, and the total time taken to select S *and* solve the reduced problem relative to the time taken to solve the original problem directly without screening. We will normally express these metrics as the *rejection fraction* $|\bar{S}|/|\mathcal{I}|$ and the *speedup factor* $t_{\text{solve}}/(t_{\text{screen}} + t^{\text{r}}_{\text{solve}})$. Here t_{solve} is the time to solve the original lasso problem, t_{screen} is the time to select the partition (screen the dictionary), and $t^{\text{r}}_{\text{solve}}$ is the time to solve the reduced lasso problem.

Not surprisingly, if we know the dual solution $\hat{\boldsymbol{\theta}}$, then it is easy to come up with a suitable partition. To see this, consider the lasso problem. For any partition $S \cup \bar{S}$, let $\hat{\mathbf{w}}$ and $\hat{\mathbf{z}}$ denote solutions of the original and reduced lasso problems, respectively. It is clear that the following always holds:

$$\begin{aligned} &{}^1\!/_2\|\mathbf{y} - B\hat{\mathbf{w}}\|_2^2 + \lambda\|\hat{\mathbf{w}}\|_1 \\ &\leq {}^1\!/_2\|\mathbf{y} - B\hat{\mathbf{z}}^{\uparrow S}\|_2^2 + \lambda\|\hat{\mathbf{z}}^{\uparrow S}\|_1 \\ &= {}^1\!/_2\|\mathbf{y} - B_{\downarrow S}\hat{\mathbf{z}}\|_2^2 + \lambda\|\hat{\mathbf{z}}\|_1 \\ &\leq {}^1\!/_2\|\mathbf{y} - B_{\downarrow S}\hat{\mathbf{w}}_{\downarrow S}\|_2^2 + \lambda\|\hat{\mathbf{w}}_{\downarrow S}\|_1. \end{aligned} \tag{11}$$

Now assume the dual solution $\hat{\boldsymbol{\theta}}$ is known, let $A(\hat{\boldsymbol{\theta}}) = \{i\colon |\hat{\boldsymbol{\theta}}^T\mathbf{b}_i| = 1\}$ denote the active constraints at $\hat{\boldsymbol{\theta}}$, and consider the particular partition $A(\hat{\boldsymbol{\theta}}) \cup \bar{A}(\hat{\boldsymbol{\theta}})$. Equation (4) shows that if $|\hat{\boldsymbol{\theta}}^T\mathbf{b}_i| < 1$ (equivalently, $i \in \bar{A}(\hat{\boldsymbol{\theta}})$), then $\hat{w}_i = 0$. Hence for this partition, $B\hat{\mathbf{w}} = B_{\downarrow A(\hat{\boldsymbol{\theta}})}\hat{\mathbf{w}}_{\downarrow A(\hat{\boldsymbol{\theta}})}$, $\|\hat{\mathbf{w}}_{\downarrow A(\hat{\boldsymbol{\theta}})}\|_1 = \|\hat{\mathbf{w}}\|_1$, and

$$\begin{aligned} &{}^1\!/_2\|\mathbf{y} - B\hat{\mathbf{w}}\|_2^2 + \lambda\|\hat{\mathbf{w}}\|_1 \\ &= {}^1\!/_2\|\mathbf{y} - B_{\downarrow A(\hat{\boldsymbol{\theta}})}\hat{\mathbf{w}}_{\downarrow A(\hat{\boldsymbol{\theta}})}\|_2^2 + \lambda\|\hat{\mathbf{w}}_{\downarrow A(\hat{\boldsymbol{\theta}})}\|_1. \end{aligned} \tag{12}$$

Equation (12) implies that for this partition the two inequalities in (11) must be equalities. It follows that $\hat{\mathbf{w}}_{\downarrow A(\hat{\boldsymbol{\theta}})}$ solves the reduced problem and $\hat{\mathbf{z}}^{\uparrow S}$ solves the original problem. Although a simple observation, this is worth stating as a theorem.

Theorem 1. *Let the solution $\hat{\boldsymbol{\theta}}$ of* (3) *(resp.* (5)*) have active set $A(\hat{\boldsymbol{\theta}})$. If $\hat{\mathbf{z}}$ is a solution of* (1) *(resp.* (2)*) with dictionary $B_{\downarrow A(\hat{\boldsymbol{\theta}})}$, then $\hat{\mathbf{z}}^{\uparrow A(\hat{\boldsymbol{\theta}})}$ solves* (1) *(resp.* (2)*). Moreover, every solution of* (1) *(resp.* (2)*) can be expressed in this way.*

The fundamental partition of $\mathcal{I}$ into $A(\hat{\boldsymbol{\theta}})$ and $\bar{A}(\hat{\boldsymbol{\theta}})$ is conceptually very important but obviously impractical. If we know $\hat{\boldsymbol{\theta}}$, then we can easily solve the primal problem (see (10)) and this makes screening and problem reduction unnecessary. As a first step towards finding a practical way to partition the features, we note that if $A(\hat{\boldsymbol{\theta}}) \subseteq S$ (screening keeps more) or equivalently $\bar{S} \subseteq \bar{A}(\hat{\boldsymbol{\theta}})$ (screening rejects less), then equation (12) holds with S replacing $A(\hat{\boldsymbol{\theta}})$. This implies that the two inequalities in (11) hold with equality for this partition, and hence we have the following corollary of Theorem 1.

Corollary 1. *Let the solution $\hat{\boldsymbol{\theta}}$ of* (3) *(resp.* (5)*) have active set $A(\hat{\boldsymbol{\theta}})$. Let $A(\hat{\boldsymbol{\theta}}) \subseteq S \subseteq \mathcal{I}$. If $\hat{\mathbf{z}}$ is a solution of* (1) *(resp.* (2)*) with dictionary $B_{\downarrow S}$, then $\hat{\mathbf{z}}^{\uparrow S}$ is a solution of* (1) *(resp.* (2)*). Moreover, every solution of* (1) *(resp.* (2)*) can be expressed in this way.*

4 Region Tests

The core idea for creating a partition of the dictionary that conforms with Corollary 1 is to bound $\hat{\boldsymbol{\theta}}$ within a compact region $\mathcal{R}$. For each feature $\mathbf{b}$, we then compute $\mu_{\mathcal{R}}(\mathbf{b}) = \max_{\boldsymbol{\theta}\in\mathcal{R}} \boldsymbol{\theta}^T\mathbf{b}$, and use this quantity to partition B [22].

We first illustrate this for the nonnegative lasso. For a compact set $\mathcal{R}$, if $\mathcal{R} = \emptyset$, all features are rejected; otherwise for each feature $\mathbf{b}_i$, $\mu_{\mathcal{R}}(\mathbf{b}_i) = \max_{\boldsymbol{\theta}\in\mathcal{R}} \boldsymbol{\theta}^T\mathbf{b}_i$ exists. Then define the partition:

$$\mathbf{b}_i \in \begin{cases} \bar{S}, & \text{if } \mu_{\mathcal{R}}(\mathbf{b}_i) < 1; \\ S, & \text{otherwise.} \end{cases} \tag{13}$$

The logic is that if $\hat{\boldsymbol{\theta}} \in \mathcal{R}$ and $\mu_{\mathcal{R}}(\mathbf{b}_i) < 1$, then $\hat{\boldsymbol{\theta}}^T\mathbf{b}_i < 1$ and hence $i \in \bar{A}(\hat{\boldsymbol{\theta}})$. Thus $\bar{S} \subseteq \bar{A}(\hat{\boldsymbol{\theta}})$, as desired.

For the lasso problem, $\hat{\boldsymbol{\theta}} \in \mathcal{R}$ and $\mu_{\mathcal{R}}(\mathbf{b}_i) < 1$ ensure $\hat{\boldsymbol{\theta}}^T\mathbf{b}_i < 1$. But in this case we also need $-1 < \hat{\boldsymbol{\theta}}^T\mathbf{b}_i$ or equivalently $\hat{\boldsymbol{\theta}}^T(-\mathbf{b}_i) < 1$. This holds if $\mu_{\mathcal{R}}(-\mathbf{b}_i) < 1$. Effectively, we must test both $\mathbf{b}_i$ and $-\mathbf{b}_i$ to account for the positive or negative sign of w_i. So for the lasso the partition is:

$$\mathbf{b}_i \in \begin{cases} \bar{S}, & \text{if } \max\{\mu_{\mathcal{R}}(\mathbf{b}_i), \mu_{\mathcal{R}}(-\mathbf{b}_i)\} < 1; \\ S, & \text{otherwise.} \end{cases} \tag{14}$$

For example, when $\mathcal{R} = \{\hat{\boldsymbol{\theta}}\}$, $i \in \bar{S}$ if: (a) $\hat{\boldsymbol{\theta}}^T\mathbf{b}_i < 1$ (nonnegative lasso) and (b) $|\hat{\boldsymbol{\theta}}^T\mathbf{b}_i| < 1$ (lasso). So $\mathcal{R} = \{\hat{\boldsymbol{\theta}}\}$, yields the ideal partition $A(\hat{\boldsymbol{\theta}}) \cup \bar{A}(\hat{\boldsymbol{\theta}})$.

A compact region bounding $\hat{\boldsymbol{\theta}}$ is thus a potentially practical means of implementing dictionary screening. From the above constructions, we see that $\hat{\boldsymbol{\theta}} \in \mathcal{R}$ ensures that the partitions (13) and (14) satisfy $\bar{S} \subseteq \bar{A}(\hat{\boldsymbol{\theta}})$. Hence the assumptions of Corollary 1 are satisfied. This is summarized in the following corollary.

Corollary 2. *Let $\mathcal{R}$ be a compact region with $\hat{\boldsymbol{\theta}} \in \mathcal{R}$. Then $\mathcal{R}$ defines a partition $S \cup \bar{S}$ of the dictionary with $\bar{S} \subseteq \bar{A}(\hat{\boldsymbol{\theta}})$.*

It will be convenient to encode the partition induced by a bounding region $\mathcal{R}$ as a rejection test $T_{\mathcal{R}}$ with $T_{\mathcal{R}}(\mathbf{b}) = 1$ if $\mathbf{b} \in \bar{S}$ (rejected) and $T_{\mathcal{R}}(\mathbf{b}) = 0$ if $\mathbf{b} \in S$ (selected). For example, the rejection test corresponding to (14) is:

$$T_{\mathcal{R}}(\mathbf{b}_i) \in \begin{cases} 1, & \text{if } \max\{\mu_{\mathcal{R}}(\mathbf{b}_i), \mu_{\mathcal{R}}(-\mathbf{b}_i)\} < 1; \\ 0, & \text{otherwise.} \end{cases} \tag{15}$$

We end this section by noting that for a given dictionary B, the partial order of subsets of features induces a partial order on screening tests. Test T' is *weaker* than test T, denoted $T' \preceq T$, if the set of features rejected by T' is a subset of the features rejected by T. For example, if $\hat{\boldsymbol{\theta}} \in \mathcal{R}$, then $T_{\mathcal{R}} \preceq T_{\{\hat{\boldsymbol{\theta}}\}}$. This is a special case of the following lemma.

Lemma 1. *If $\mathcal{R}_1 \subseteq \mathcal{R}_2$, then $T_{\mathcal{R}_2} \preceq T_{\mathcal{R}_1}$.*

By Lemma 1, the region test for $\mathcal{R}_1 \subset \mathcal{R}_2$ can potentially reject more features than the test for $\mathcal{R}_2$.

4.1 The Sphere-Hyperplane Region Architecture

We now consider particular forms of bounding regions for $\hat{\boldsymbol{\theta}}$. A natural form of bounding region consists of the intersection of a spherical bound with a finite number of half spaces. The spherical bound arises naturally once we know a dual feasible point, and half spaces arise naturally since these define the dual feasible region $\mathcal{F}$ (see (3)), and are integral to the projection of a point onto $\mathcal{F}$ (see (7)).

The intersection of a closed spherical ball $S(\mathbf{q}, r) = \{\mathbf{z} : \|\mathbf{z} - \mathbf{q}\|_2 \leq r\}$ with center $\mathbf{q}$ and radius r, and m half spaces $\mathbf{n}_i^T\boldsymbol{\theta} \leq c_i$, $i = 1, \ldots, m$, gives the region:

$$\mathcal{R} = \{\boldsymbol{\theta} \colon \|\boldsymbol{\theta} - \mathbf{q}\|_2 \leq r\} \cap \cap_{i=1}^m \{\boldsymbol{\theta} \colon \mathbf{n}_i^T\boldsymbol{\theta} \leq c_i\} .$$

To form the corresponding region test, we find $\mu(\mathbf{b}) = \max_{\boldsymbol{\theta}\in\mathcal{R}} \boldsymbol{\theta}^T\mathbf{b}$ by solving the optimization problem:

$$\begin{aligned} \min_{\boldsymbol{\theta}} \quad & (-\boldsymbol{\theta}^T\mathbf{b}) \\ \text{s.t.} \quad & (\boldsymbol{\theta} - \mathbf{q})^T(\boldsymbol{\theta} - \mathbf{q}) - r^2 \leq 0 \\ & \mathbf{n}_i^T\boldsymbol{\theta} - c_i \leq 0, \quad i = 1, \ldots, m. \end{aligned} \tag{16}$$

Once $\mu(\mathbf{b})$ is known, (15) gives the corresponding screening test.

Using the change of variable $\mathbf{z} = (\boldsymbol{\theta} - \mathbf{q})/r$, problem (16) can be simplified to:

$$\begin{aligned} \bar{\mu}(\mathbf{b}) = \min_{\mathbf{z}} \quad & (-\mathbf{z}^T\mathbf{b}) \\ \text{s.t.} \quad & \mathbf{z}^T\mathbf{z} - 1 \leq 0 \\ & \mathbf{n}_i^T\mathbf{z} + \psi_i \leq 0, \quad i = 1, \ldots, m. \end{aligned} \tag{17}$$

where $\psi_i = (\mathbf{n}_i^T\mathbf{q} - c_i)/r$. The solution of (16) is then given by $\mu(\mathbf{b}) = \mathbf{q}^T\mathbf{b} + r\bar{\mu}(\mathbf{b})$. The solution of (17) only depends on the projection of $\mathbf{b}$ onto the subspace spanned by $\mathbf{n}_i$, $i = 1, \ldots, m$. Problem (17) is thus a constrained optimization problem in $\mathbb{R}^m$.

Increasing m results in tests with the potential to reject more features, but which are also more complex and time consuming to execute. In the following two subsections, we discuss the simplest cases: $m = 0$ (sphere tests), and $m = 1$ (dome tests). This gives insight into basic tests and makes connections with the literature.

4.2 Sphere Tests

Consider bounding $\hat{\boldsymbol{\theta}}$ within a closed spherical ball $S(\mathbf{q}, r) = \{\mathbf{z} : \|\mathbf{z} - \mathbf{q}\|_2 \leq r\}$ with center $\mathbf{q}$ and radius r. This bound gives a simple, efficiently implemented

test, and it is also a useful building block for more complex tests. We first determine a close form expression for $\mu_{S(\mathbf{q},r)}(\mathbf{b}) = \max_{\boldsymbol{\theta}\in S(\mathbf{q},r)} \boldsymbol{\theta}^T\mathbf{b}$. An expression for a *sphere test* $T_{S(\mathbf{q},r)}$ then follows from (15).

Lemma 2. *For $S(\mathbf{q},r) = \{\mathbf{z} : \|\mathbf{z}-\mathbf{q}\|_2 \le r\}$ and $\mathbf{b} \in \mathbb{R}^n$:*

$$\mu_{S(\mathbf{q},r)}(\mathbf{b}) = \mathbf{q}^T\mathbf{b} + r\|\mathbf{b}\|_2. \tag{18}$$

Theorem 2. *The screening test for the sphere $S(\mathbf{q},r)$ is:*

$$T_{S(\mathbf{q},r)}(\mathbf{b}) = \begin{cases} 1, \text{ if } V_l(\|\mathbf{b}\|_2) < \mathbf{q}^T\mathbf{b} < V_u(\|\mathbf{b}\|_2); \\ 0, \text{ otherwise.} \end{cases} \tag{19}$$

where $V_u(t) = 1 - rt$ and for the lasso $V_l(t) = -V_u(t)$, and for the nonnegative lasso $V_l(t) = -\infty$.

For the lasso problem, the test (19) can also be written in the form:

$$T_{S(\mathbf{q},r)}(\mathbf{b}) = \begin{cases} 1, & \text{if } |\mathbf{q}^T\mathbf{b}| < 1 - r\|\mathbf{b}\|_2; \\ 0, & \text{otherwise.} \end{cases} \tag{20}$$

Theorem 2 defines a parametric family of tests: $\{\mathrm{ST}(\mathbf{q},r)\colon \mathbf{q} \in \mathbb{R}^n, r \ge 0\}$, where $\mathrm{ST}(\mathbf{q},r)$ denotes the sphere test with center $\mathbf{q}$ and radius r. To use a sphere test one first selects values of $\mathbf{q}$ and r so that $S(\mathbf{q},r)$ bounds $\hat{\boldsymbol{\theta}}$. We call this the *parameter selection problem*. By Lemma 1, a tighter bound has potential for better screening. So using only the information provided, and limited computation, we want to select $\mathbf{q}$ and r to give the "best bound". This is a design problem involving a trade-off between the computation cost to select $\mathbf{q}$ and r and the resultant screening performance. Hence we don't expect there is a "best answer". We outline below several selection methods.

4.2.1 Parameter selection

If we know a dual feasible point $\boldsymbol{\theta}_F \in \mathcal{F}$, then $\hat{\boldsymbol{\theta}}$ can't be further away from $\mathbf{y}/\lambda$ than $\boldsymbol{\theta}_F$. This gives the basic spherical bound:

$$\|\hat{\boldsymbol{\theta}} - \mathbf{y}/\lambda\|_2 \le \|\boldsymbol{\theta}_F - \mathbf{y}/\lambda\|_2, \tag{21}$$

with center $\mathbf{q} = \mathbf{y}/\lambda$ and radius $r = \|\boldsymbol{\theta}_F - \mathbf{y}/\lambda\|_2$. In particular, the dual solution $\hat{\boldsymbol{\theta}}(\lambda_{\max}) = \mathbf{y}/\lambda_{\max}$. This gives a particular instance of (21):

$$\|\hat{\boldsymbol{\theta}} - \mathbf{y}/\lambda\|_2 \le |1/\lambda - 1/\lambda_{\max}|\ \|\mathbf{y}\|_2. \tag{22}$$

This bound is shown in Fig. 1(c) as the larger sphere in solid red. The bound (22) requires only the specification of the lasso problem and the computation of $\lambda_{\max}$. We call it the *default spherical bound*.

Better bounds are possible with additional computation or if additional information is supplied. For example, [27] observed that to obtain a feasible point $\boldsymbol{\theta}_F$ closer to $\hat{\boldsymbol{\theta}}$ than $\mathbf{y}/\lambda_{\max}$ one can first run K steps of the homotopy algorithm on (1). This gives the solution $\hat{\mathbf{w}}_K$ of the instance $(\mathbf{y}, \lambda_K)$, $\lambda_K > \lambda$, for the K-th breakpoint on the (primal) regularization path. Effectively, this first solves the lasso problem for $\lambda_K > \lambda$, and then uses this solution to help screen for the actual instance to be solved. The sphere center can also be moved away from $\mathbf{y}/\lambda$. Examples include the sphere tests ST2 and ST3 in [23] derived in the setting of unit norm $\mathbf{y}$ and $\mathbf{b}_i$. In addition, [27] noted that if the dual solution $\hat{\boldsymbol{\theta}}_0$ is known for an instance $(\mathbf{x}, \lambda_0)$, then $\|\hat{\boldsymbol{\theta}}(\lambda) - \hat{\boldsymbol{\theta}}_0\|_2 \le |1/\lambda - 1/\lambda_0|\ \|\mathbf{y}\|_2$ (this is discussed further below). This leverages a solved instance to give a spherical bound centered at $\mathbf{q} = \hat{\boldsymbol{\theta}}_0$.

4.2.2 Connections with the Literature

A variety of existing screening tests for the lasso are sphere tests. We give a quick synopsis that places these tests within the context of our exposition. The Basic SAFE-LASSO test [34] and the test ST1 in [23, Sect. 2] are sphere tests based on the default spherical bound (22). The SAFE-LASSO test [34, Theorem 2] is also a sphere test. It assumes a dual feasible point $\boldsymbol{\theta}_0$ is given and uses this to improve the default spherical bound centered at $\mathbf{y}/\lambda$. The sphere tests ST2 and ST3 in [23, Sect. 2] use spherical bounds not centered at $\mathbf{y}/\lambda$. We will comment further on the test ST3 at the end of §4.6. The core test used in [27] is a sphere test with center $\hat{\boldsymbol{\theta}}_0 = \hat{\boldsymbol{\theta}}(\lambda_0)$, where $\hat{\boldsymbol{\theta}}_0(\lambda_0)$ is the dual solution at λ_0, and radius $|1/\lambda - 1/\lambda_0|\ \|\mathbf{y}\|_2$. This bound follows from the nonexpansive property of projection onto a convex set:

$$\begin{aligned} \|\hat{\boldsymbol{\theta}}(\lambda) - \hat{\boldsymbol{\theta}}(\lambda_0)\|_2 &\le \|\mathbf{y}/\lambda - \mathbf{y}/\lambda_0\|_2 \\ &= |1/\lambda - 1/\lambda_0|\ \|\mathbf{y}\|_2. \end{aligned} \tag{23}$$

The Strong Rule [20] is also a sphere test for the lasso problem. For notational simplicity, let the features and the target vector $\mathbf{y}$ have unit norm. The Strong Rule discards feature $\mathbf{b}_i$ if $|\mathbf{b}_i^T\mathbf{y}| < 2\lambda - \lambda_{\max}$. This is a sphere test with center $\mathbf{q} = \mathbf{y}/\lambda$ and radius $r_{sr} = (\lambda_{\max} - \lambda)/\lambda$. However, $\hat{\boldsymbol{\theta}}$ may not be bounded within this sphere. So the Strong Rule can yield false rejections. To see why, note that $\hat{\boldsymbol{\theta}}$ is bounded within the default sphere (center $\mathbf{y}/\lambda$, radius $r = 1/\lambda - 1/\lambda_{\max}$). The Strong Rule is using a sphere centered at the same location but with only a fraction of this radius: $r_{sr} = r\lambda_{\max}$. A more advanced version of the Strong Rule, the Strong Sequential Rule [20], assumes a solution $\hat{\mathbf{w}}_0$ of the lasso instance $(\mathbf{y}, \lambda_0)$ is available, where $\lambda_0 > \lambda$. It then forms the residual $\mathbf{r}_0 = \mathbf{y} - B\hat{\mathbf{w}}_0$ and screens the lasso instance $(\mathbf{y}, \lambda)$ using the test $|\mathbf{b}_i^T\mathbf{r}_0| < 2\lambda - \lambda_0$. This is also a sphere test. To see this, use (4) to write $\mathbf{r}_0 = \mathbf{y} - B\hat{\mathbf{w}} = \lambda_0\hat{\boldsymbol{\theta}}_0$. Then the test becomes $|\mathbf{b}_i^T\hat{\boldsymbol{\theta}}_0| < 1 - r_{ssr}$ with $r_{ssr} = (1/\lambda - 1/\lambda_0)2\lambda$. This is a sphere test with center $\hat{\boldsymbol{\theta}}_0$ and radius $r_{ssr} = 2\lambda r$ where r is the radius of the known bounding sphere (23). When $\lambda < 0.5$, this test may also yield false rejections.

The SIS test in [19] is framed in a probabilistic setting and is not intended for lasso screening. Nevertheless, if we translate SIS into our setting it is a sphere test for a lasso problem with appropriately selected λ.

SIS assumes a dictionary $B \in \mathbb{R}^{n\times p}$ of standardized vectors (features of unit norm) and computes the vector of (marginal) correlations $\boldsymbol{\rho} = B^T\mathbf{y}$. Then given $0 < \gamma < 1$, it selects the top $[\gamma n]$ features ranked by $|\rho_i|$. Assume for simplicity that the values $|\rho_i|$ are distinct and let t_γ denote the value of $|\rho_i|$ for the $[\gamma n]$-th feature in the ranking. The SIS rejection criterion can then be written as $|\mathbf{b}_i^T\mathbf{y}| < t_\gamma$. We now form a lasso problem with dictionary B, target vector $\mathbf{y}$, and a value $\lambda/\lambda_{\max}$ to be decided. For simplicity of notation, assume that $\mathbf{y}$ has unit norm. Then the default spherical bound for the dual solution of the lasso has center $\mathbf{y}/\lambda$ and radius $r = 1/\lambda - 1/\lambda_{\max}$, and the corresponding sphere test is $|\mathbf{b}_i^T\mathbf{y}| < \lambda(1-r)$. Equating the right hand sides of the above test expressions, and using some algebra shows that if we take $\lambda/\lambda_{\max} = (1+t_\gamma)/(1+\lambda_{\max}) < 1$, then SIS is the default sphere test for this particular lasso problem.

4.3 Sphere Plus Halfspace Tests

Now consider a region test based on the nonempty intersection of a spherical ball $\{\mathbf{z}\colon \|\mathbf{z}-\mathbf{q}\|_2 \le r\}$ and one closed half space $\{\mathbf{z}\colon \mathbf{n}^T\mathbf{z} \le c\}$. Here $\mathbf{n}$ is the unit normal to the half space and $c \ge 0$. This yields the region $D(\mathbf{q}, r; \mathbf{n}, c) = \{\mathbf{z}\colon \mathbf{n}^T\mathbf{z} \le c, \|\mathbf{z}-\mathbf{q}\|_2 \le r\}$. This illustrated in Fig. 2(a). For brevity we refer to this form of region as a dome.

The following features of the dome $D(\mathbf{q}, r; \mathbf{n}, c)$ will be useful. We call the point $\mathbf{q}_d$ on the bounding hyperplane and the line passing through $\mathbf{q}$ in the direction of the hyperplane normal the *dome center*. The signed distance from $\mathbf{q}$ to $\mathbf{q}_d$ in the direction $-\mathbf{n}$ is a fraction ψ_d of the radius r of the sphere. We call the maximum straight line distance r_d one can move from $\mathbf{q}_d$ within the dome and hyperplane the *dome radius*. Under the sign convention indicated above, simple Euclidean geometry gives the following relationships:

$$\psi_d = (\mathbf{n}^T\mathbf{q} - c)/r, \tag{24}$$

$$\mathbf{q}_d = \mathbf{q} - \psi_d r\mathbf{n}, \tag{25}$$

$$r_d = r\sqrt{1-\psi_d^2}. \tag{26}$$

To ensure that the dome is nondegenerate (a nonempty and proper subset of each region), we need $\mathbf{q}_d$ to be inside the sphere. Hence we require $-1 \le \psi_d \le 1$. So we need $\mathbf{q}^T\mathbf{n} \ge c - r$, this ensures that the intersection is a proper subset of the sphere and the half space; and we need $\mathbf{q}^T\mathbf{n} \le c + r$, this ensures the intersection is nonempty.

To find $\mu(\mathbf{b}) = \max_{\boldsymbol{\theta}\in D(\mathbf{q},r;\mathbf{n},c)} \boldsymbol{\theta}^T\mathbf{b}$, for $\mathbf{b} \in \mathbb{R}^n$, we solve the optimization problem (16) with $m = 1$. Particular instances of this problem were solved in [35, Appendix A] (by solving a Lagrange dual problem) and in [24, §3], [25] (by directly solving a primal problem). Both approaches can be extended to solve the general problem (16) with $m = 1$. This yields the following lemma, and the dome screening test.

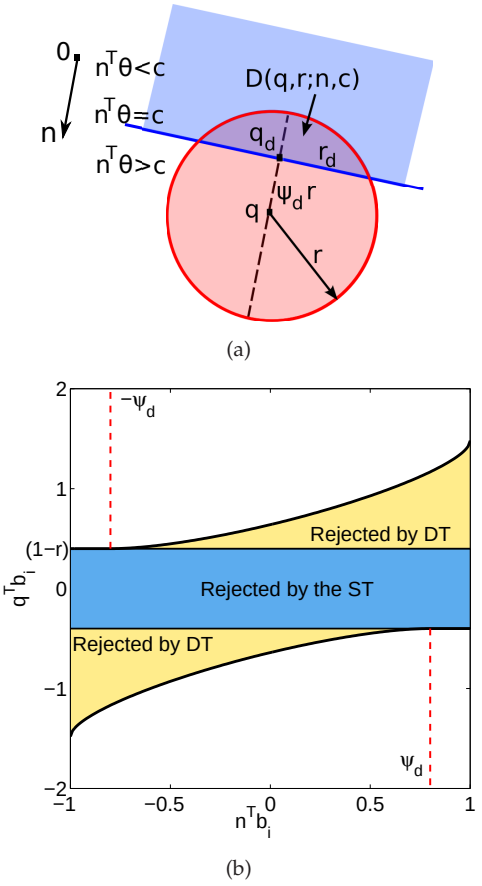

Fig. 2. **(a)** A general dome region $D(\mathbf{q}, r; \mathbf{n}, c)$ shown for $0 < \psi_d < 1$ and the dome consisting of less than half the sphere. **(b)** The rejection area (shaded) of a lasso dome test.

Lemma 3. *Fix a dome $D = D(\mathbf{q}, r; \mathbf{n}, c)$ with $|\psi_d| \le 1$. Then for $\mathbf{b} \in \mathbb{R}^n$,*

$$\mu_D(\mathbf{b}) = \mathbf{q}^T\mathbf{b} + M_1(\mathbf{n}^T\mathbf{b}, \|\mathbf{b}\|_2),$$

where $M_1(t_1, t_2)$ is the function

$$M_1(t_1,t_2) = \begin{cases} rt_2, & \text{if } t_1 < -\psi_d t_2; \\ -\psi_d r t_1 + r\sqrt{t_2^2 - t_1^2}\sqrt{1-\psi_d^2}, & \text{if } t_1 \ge -\psi_d t_2. \end{cases} \tag{27}$$

Theorem 3. *The screening test for a nondegenerate dome $D(\mathbf{q}, r; \mathbf{n}, c)$ is:*

$$T_{D(\mathbf{q},r;\mathbf{n},c)}(\mathbf{b}) = \begin{cases} 1, & \text{if } V_l(\mathbf{n}^T\mathbf{b}, \|\mathbf{b}\|_2) < \mathbf{q}^T\mathbf{b} < V_u(\mathbf{n}^T\mathbf{b}, \|\mathbf{b}\|_2); \\ 0, & \text{otherwise;} \end{cases} \tag{28}$$

where $V_u(t_1, t_2) = 1 - M_1(t_1, t_2)$ and for the lasso $V_l(t_1, t_2) = -V_u(-t_1, t_2)$, and for the nonnegative lasso $V_l(t_1, t_2) = -\infty$.

We denote a dome test by $\mathrm{DT}(\mathbf{q}, r; \mathbf{n}, c)$. Although defined piecewise, the functions V_u and V_l in Theorem 3 are continuous and smooth: $V_u, V_l \in C^1$. This can be checked using simple calculus. The parameters r and c of the dome do not appear as arguments in the test but play a role through M_1. The test simplifies for unit

norm features. In that case, $t_2 = \|\mathbf{b}_i\|_2 = 1$ and M_1, V_u and V_l are only functions of t_1.

To gain some insight into this test, consider the situation when $r < 1$ and all features have unit norm. We can factor the test into the composition of two functions: a linear map $\mathbf{b}_i \mapsto [\mathbf{q},\ \mathbf{n}]^T\mathbf{b}_i$ and a two-dimensional decision function H_{r,ψ_d}:

$$H_{r,\psi_d}(s,t) = \begin{cases} 1, & \text{if } V_l(t) < s < V_u(t); \\ 0, & \text{otherwise;} \end{cases}$$

where $s = \mathbf{q}^T\mathbf{b}_i \in [-\|\mathbf{q}\|_2, \|\mathbf{q}\|_2]$, $t = \mathbf{n}^T\mathbf{b}_i \in [-1, 1]$, and $V_u(t)$, $V_l(t)$ are given in Theorem 3 with $t = t_1$ and $t_2 = 1$. We can display the test rejection region by plotting $V_l(t)$ and $V_u(t)$ versus t as shown in Fig. 2(b). For the lasso, the rejection region has upper and lower boundaries. The sections of the boundaries with $V_u(t) = (1 - r)$ and $V_l(t) = -(1 - r)$, correspond to the sphere test $T_{S(\mathbf{q},r)}$. If feature $\mathbf{b}_i$ maps into the shaded region in the figure, then $\mathbf{b}_i$ is rejected. The lightly shaded (yellow) area indicates the extra rejection power of the dome test over the underlying sphere test. For a given value of $\mathbf{q}^T\mathbf{b}_i > 0$, the dome test lowers the bar for rejection as $\mathbf{n}^T\mathbf{b}_i$ increases.

4.3.1 Parameter selection

Now consider the parameter selection problem. Since we have discussed parameter selection for a spherical bound, we assume $S(\mathbf{q}, r)$ is given and give examples of bounding $\hat{\boldsymbol{\theta}}$ within a suitable half space.

Each constraint of the dual problem $\mathbf{b}^T\boldsymbol{\theta} \le 1$ bounds $\hat{\boldsymbol{\theta}}$. This half space has $\mathbf{n} = \mathbf{b}/\|\mathbf{b}\|_2$ and $c = 1/\|\mathbf{b}\|_2$. The resultant dome is nonempty since both the sphere and the half space contain $\hat{\boldsymbol{\theta}}$. To ensure it is proper, we require $\mathbf{q}^T\mathbf{b} \ge 1 - r\|\mathbf{b}\|_2$. This means that the sphere test does not reject the feature $\mathbf{b}$. In particular, we can select $\mathbf{b}$ to minimize the disk radius r_d. To do so, we maximize ψ_d given by (24):

$$\mathbf{b}_g = \arg\max_{\mathbf{b}\in\mathfrak{B}} \frac{\mathbf{b}^T\mathbf{q} - 1}{\|\mathbf{b}\|_2} \tag{29}$$

For unit norm features, (29) selects the feature most correlated with $\mathbf{q}$. If in addition, $\mathbf{q} = \mathbf{y}/\lambda$, then (29) yields $\mathbf{b}_g = \mathbf{b}_{\max}$. Selecting the default spherical bound and using (29) gives the specific dome:

$$D(\mathbf{y}/\lambda, |1/\lambda - 1/\lambda_{\max}|\ \|\mathbf{y}\|_2; \mathbf{b}_g/\|\mathbf{b}_g\|_2, 1/\|\mathbf{b}_g\|_2). \tag{30}$$

We call this the *default dome bound*. When $\mathbf{y}$ and all the features have unit norm, this simplifies to $D(\mathbf{y}/\lambda, |1/\lambda - 1/\lambda_{\max}|; \mathbf{b}_{\max}, 1)$. This dome is illustrated in Fig. 3(left).

If $\hat{\boldsymbol{\theta}}_0$ is the dual solution of an instance $(\mathbf{y}_0, \lambda_0)$, then $\hat{\boldsymbol{\theta}}_0$ lies on the boundary of $\mathcal{F}$. Moreover, its *optimality* for $(\mathbf{y}_0, \lambda_0)$ ensures that it satisfies the inequalities (7). Hence for each $\boldsymbol{\theta} \in \mathcal{F}$:

$$(\mathbf{y}_0/\lambda_0 - \hat{\boldsymbol{\theta}}_0)^T\boldsymbol{\theta} \le (\mathbf{y}_0/\lambda_0 - \hat{\boldsymbol{\theta}}_0)^T\hat{\boldsymbol{\theta}}_0. \tag{31}$$

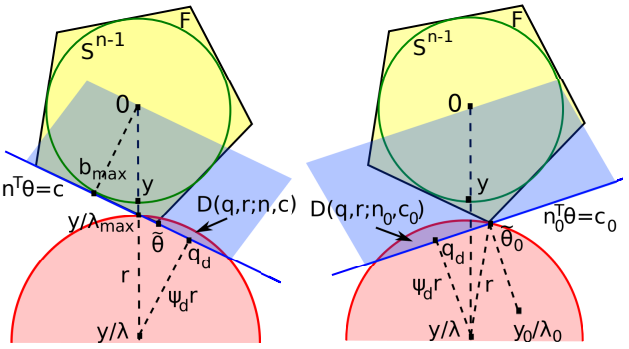

Fig. 3. Two dome tests for unit norm features and target vector. Left: The dome (30) based on the feasible point $\mathbf{y}/\lambda_{\max}$ and the closed half space $\mathbf{b}_{\max}^T\boldsymbol{\theta} \le 1$. Right: The dome (33) based on a solved instance $(\mathbf{y}_0, \lambda_0, \hat{\boldsymbol{\theta}}_0)$.

Since $0 \in \mathcal{F}$, the right hand side is nonnegative. Therefore this inequality bounds $\mathcal{F}$ in the closed half space $\mathbf{n}_0^T\boldsymbol{\theta} \le c_0$ with

$$\begin{aligned} r_0 &= \|\mathbf{y}_0/\lambda_0 - \hat{\boldsymbol{\theta}}_0\|_2, \\ \mathbf{n}_0 &= (\mathbf{y}_0/\lambda_0 - \hat{\boldsymbol{\theta}}_0)/r_0, \\ c_0 &= \mathbf{n}_0^T\hat{\boldsymbol{\theta}}_0. \end{aligned} \tag{32}$$

The intersection of this half space with the bounding sphere $S(\mathbf{q}, r)$ is nonempty and it is proper if $\psi_d \ge -1$. To check this condition we note that

$$\begin{aligned} \psi_d &= \frac{\mathbf{n}_0^T\mathbf{q} - \mathbf{n}_0^T\hat{\boldsymbol{\theta}}_0}{r} \\ &= \frac{(\mathbf{y}_0/\lambda_0 - \hat{\boldsymbol{\theta}}_0)^T}{r_0}\frac{(\mathbf{q} - \hat{\boldsymbol{\theta}}_0)}{\|\mathbf{q} - \hat{\boldsymbol{\theta}}_0\|_2}\frac{\|\mathbf{q} - \hat{\boldsymbol{\theta}}_0\|_2}{r} \\ &= \cos\beta\frac{\|\mathbf{q} - \hat{\boldsymbol{\theta}}_0\|_2}{r} \end{aligned}$$

where β is the angle between n_0 and $\mathbf{q} - \hat{\boldsymbol{\theta}}_0$. So if $\cos\beta > 0$ or $\hat{\boldsymbol{\theta}}_0 \in S(\mathbf{q}, r)$, then the dome is proper. For example, $\mathbf{q} = \mathbf{y}/\lambda$ and $r = \|\mathbf{y}/\lambda - \hat{\boldsymbol{\theta}}_0\|_2$, yields the proper dome:

$$D(\mathbf{y}/\lambda, \|\mathbf{y}/\lambda - \hat{\boldsymbol{\theta}}_0\|_2; \mathbf{n}_0, c_0). \tag{33}$$

This dome illustrated in Fig. 3(right).

4.3.2 Connections with the Literature

Specific dome tests were introduced in [35, §2.4] and [24, §3]. The dome test discussed in [24] is based on the default dome bound (30) for unit norm features and unit norm $\mathbf{y}$.

The SAFE-LASSO test in [35, §2.4] is a dome test specifically designed for screening and solving lasso problems at points along the regularization path. A triple $(\mathbf{y}, \lambda_0, \hat{\boldsymbol{\theta}}_0)$ is assumed given where $\hat{\boldsymbol{\theta}}_0$ is the dual solution for the instance $(\mathbf{y}, \lambda_0)$. The test exploits this information to screen the dictionary for an instance $(\mathbf{y}, \lambda)$ with $\lambda < \lambda_0 \le \lambda_{\max}$. We show below that the dome employed is (33) with $\mathbf{y}_0 = \mathbf{y}$.

The solution provided in [35, §2.4] entails specifying a bounding sphere and a half space, then solving the

corresponding version of problem (16). The selected half space is $g^T(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_0) \geq 0$ where $g = \nabla G(\hat{\boldsymbol{\theta}}_0) = -\mathbf{y}/\lambda_0 + \hat{\boldsymbol{\theta}}_0$ is the gradient of the dual objective for the solved instance evaluated at $\hat{\boldsymbol{\theta}}_0$ (up to positive scaling). The spherical bound is obtained by scaling $\hat{\boldsymbol{\theta}}_0$ to obtain the closest feasible solution to $\mathbf{y}/\lambda$. This can be specified by letting $r_0 = \|\mathbf{y}/\lambda_0 - \hat{\boldsymbol{\theta}}_0\|_2$ and setting $\mathbf{q} = \mathbf{y}/\lambda$, $\hat{r} = \min_{s\in[-1,1]} \|s\hat{\boldsymbol{\theta}}_0 - \mathbf{y}/\lambda\|_2$, $\mathbf{n} = (\mathbf{y}/\lambda_0 - \hat{\boldsymbol{\theta}}_0)/r_0$ and $c = \mathbf{n}^T\hat{\boldsymbol{\theta}}_0$.

Assume $\lambda \leq \lambda_0 < \lambda_{\max}$ and let $\hat{s}(\lambda)$ denote the optimal value of s in the definition of $\hat{r}$. By the optimality of $\hat{\boldsymbol{\theta}}_0$ for the instance $(\mathbf{y}, \lambda_0)$, we must have $\hat{s}(\lambda_0) = 1$. In addition, it must hold that $\mathbf{y}^T\hat{\boldsymbol{\theta}}_0 \geq 0$ otherwise the feasible point $-\hat{\boldsymbol{\theta}}_0$ would be closer to $\mathbf{y}/\lambda$. By simple calculus we then determine that $\hat{s}(\lambda) = \min\{1, (\mathbf{y}^T\hat{\boldsymbol{\theta}}_0)/(\lambda\|\hat{\boldsymbol{\theta}}_0\|_2^2)\}$. It follows that for all $\lambda < \lambda_0$, $\hat{s} = 1$. Hence for $\lambda < \lambda_0$ we can take $\hat{r} = \|\mathbf{y}/\lambda - \hat{\boldsymbol{\theta}}_0\|_2$. Thus for $\lambda < \lambda_0$, SAFE-LASSO uses the dome (33) with the constraint $\mathbf{y}_0 = \mathbf{y}$.

4.4 Iteratively Refined Bounds

Under favorable circumstances, it is possible to refine a sphere $S(\mathbf{q},r)$ bounding $\hat{\boldsymbol{\theta}}$ to obtain a bounding sphere of smaller radius. Let the half space $(\mathbf{n}, c)$ also bound $\hat{\boldsymbol{\theta}}$ and its intersection with $S(\mathbf{q}, r)$ result in a dome $D = D(\mathbf{q}, r; \mathbf{n}, c)$ with parameters ψ_d, $\mathbf{q}_d$, and r_d. Since D is a bounded convex set, there exists a unique sphere of smallest radius that bounds D. This is called the *circumsphere* of D. We claim that if $0 < \psi_d \leq 1$, or equivalently $\mathbf{q} \notin D$, then the circumsphere of D is $S(\mathbf{q}_d, r_d)$. In this case, r_d is strictly smaller than r and $S(\mathbf{q}_d, r_d)$ is a tighter spherical bound on $\hat{\boldsymbol{\theta}}$. This is summarized below.

Lemma 4. *Let $S = S(\mathbf{q}, r)$ and the half space $(\mathbf{n}, c)$ bound the dual solution $\hat{\boldsymbol{\theta}}$, with the resulting dome $D = D(\mathbf{q}, r; \mathbf{n}, c)$ satisfying $0 < \psi_d \leq 1$. Then $S(\mathbf{q}_d, r_d)$ is the circumsphere of D and hence bounds $\hat{\boldsymbol{\theta}}$.*

The construction set forth in Lemma 4 can be used iteratively provided suitable half spaces can be found. The features $\mathbf{b} \in \mathfrak{B}$ provide one source of potential half spaces. At step k, we have a bounding sphere $S_k = S(\mathbf{q}_k, r_k)$ and we seek $\mathbf{b} \in \mathfrak{B}$ defining a dome with $\mathbf{n} = \mathbf{b}/\|\mathbf{b}\|_2$ and $c = 1/\|\mathbf{b}\|_2$ such that

$$0 < \psi_k = (\mathbf{q}_k^T\mathbf{n} - c)/r_k \leq 1. \tag{34}$$

If no such $\mathbf{b} \in \mathfrak{B}$ exists, then stop; otherwise set $\mathbf{n}_k = \mathbf{n}$, $c_k = c$ and

$$\mathbf{q}_{k+1} = \mathbf{q}_k - \psi_k r_k \mathbf{n}_k, \tag{35}$$

$$r_{k+1} = r_k\sqrt{1 - \psi_k^2}, \tag{36}$$

to obtain a tighter bounding sphere $S_{k+1} = S(\mathbf{q}_{k+1}, r_{k+1})$. A greedy strategy selects $\mathbf{b} \in \mathfrak{B}$ at step k to minimize r_{k+1}, or equivalently to maximize ψ_k:

$$\mathbf{b}^{(k)} = \arg\max_{\mathbf{b}\in\mathfrak{B}} \frac{\mathbf{b}^T\mathbf{q}_k - 1}{\|\mathbf{b}\|_2}\ . \tag{37}$$

When the features all have the same norm, this reduces to maximizing the inner product between $\mathbf{b}$ and $\mathbf{q}_k$. This has a simple interpretation. S_k can be thought of as a location bound on $\hat{\boldsymbol{\theta}}$ with the center $\mathbf{q}_k$ the "estimate" of $\hat{\boldsymbol{\theta}}$ given the bound. The greedy strategy selects $\mathbf{b}$ by maximizing its alignment with the current estimate $\mathbf{q}_k$ of $\hat{\boldsymbol{\theta}}$. Since $\hat{\boldsymbol{\theta}}$ is proportional to the optimal residual in the primal problem (see (4)), the greedy strategy is selecting features that "best correlate" with the current estimate of the optimal residual.

We will make use of the above construction in the following section and in §4.6.

4.5 Are More Half Spaces Worthwhile?

We have examined region tests defined by the intersection of a bounding sphere and one half space ($m = 1$), and have shown that, in general, these have additional rejection power over the simpler sphere tests ($m = 0$). Are more complex tests worthwhile? To examine this, we go one step further and examine the region test defined by the intersection of a bounding sphere and two half spaces ($m = 2$). Examining the relative performance of this test will allow us to determine where we currently stand in the trade-off between rejection power and computational efficiency.

We seek the region test based on the intersection of a sphere $S(\mathbf{q}, r) = \{\boldsymbol{\theta} : \|\boldsymbol{\theta} - \mathbf{q}\|_2 \leq r\}$ and two closed half spaces $H_i = \{\boldsymbol{\theta} : \mathbf{n}_i^T\boldsymbol{\theta} \leq c_i\}$, where $\mathbf{n}_i$ is the the unit normal to H_i and $c_i \geq 0$, $i = 1, 2$. This region is denoted by $\mathcal{R}(\mathbf{q}, r; \mathbf{n}_1, c_1; \mathbf{n}_2, c_2)$. We call the corresponding test a *Two Hyperplane Test* (THT).

Each half space $H_i = \{\boldsymbol{\theta}\colon \mathbf{n}_i^T\boldsymbol{\theta} \leq c_i\}$ intersects the sphere forming a dome with parameters $\psi_i = (\mathbf{n}_i^T\mathbf{q} - c_i)/r$, $\mathbf{q}_i = \mathbf{q} - \psi_i r\mathbf{n}_i$, and $r_i = r\sqrt{1 - \psi_i^2}$, $i = 1, 2$. To ensure each intersection $H_i \cap S(\mathbf{q}, r)$ is nonempty and proper, we need $-1 \leq \psi_i \leq 1$, $i = 1, 2$, and to ensure the two half spaces intersect within the sphere, we need $\arccos\psi_1 + \arccos\psi_2 \geq \arccos(\mathbf{n}_1^T\mathbf{n}_2)$. Under these conditions, $\mathcal{R}(\mathbf{q}, r; \mathbf{n}_1, c_1; \mathbf{n}_2, c_2)$ is a nonempty, proper subset of the sphere and each half space.

To find $\mu_{\mathcal{R}}(\mathbf{b}) = \max_{\mathbf{b}\in\mathcal{R}} \boldsymbol{\theta}^T\mathbf{b}$ we solve the optimization problem (16) with $m = 2$. Using standard techniques, this problem can be solved in closed form yielding the expressions for $\mu_{\mathcal{R}}$ in the following lemma. The corresponding test then follows from (15).

Lemma 5. *Fix the region $\mathcal{R} = \mathcal{R}(\mathbf{q}, r; \mathbf{n}_1, c_1; \mathbf{n}_2, c_2)$ and let ψ_i satisfy $|\psi_i| \leq 1$, $i = 1, 2$, and $\arccos\psi_1 + \arccos\psi_2 \geq \arccos(\mathbf{n}_1^T\mathbf{n}_2)$. Let $h(x, y, z) = \sqrt{(1-\tau^2)z^2 + 2\tau xy - x^2 - y^2}$, where $\tau = \mathbf{n}_1^T\mathbf{n}_2$. Then for $\mathbf{b} \in \mathbb{R}^n$,*

$$\mu_{\mathcal{R}}(\mathbf{b}) = \mathbf{q}^T\mathbf{b} + M_2(\mathbf{n}_1^T\mathbf{b}, \mathbf{n}_2^T\mathbf{b}, \|\mathbf{b}\|_2) \tag{38}$$

where

$$M_2(t_1,t_2,t_3) = \begin{cases} rt_3, & \text{if } (a);\\ -rt_2\psi_2 + r\sqrt{t_3^2-t_2^2}\sqrt{1-\psi_2^2}, & \text{if } (b);\\ -rt_1\psi_1 + r\sqrt{t_3^2-t_1^2}\sqrt{1-\psi_1^2}, & \text{if } (c);\\ -\frac{r}{1-\tau^2}\left[(\psi_1-\tau\psi_2)t_1+(\psi_2-\tau\psi_1)t_2\right]+ & \\ \quad \frac{r}{1-\tau^2}h(\psi_1,\psi_2,1)h(t_1,t_2,t_3), & \text{otherwise;} \end{cases}$$

and conditions $(a),(b),(c)$ *are given by*

(a) $t_1 < -\psi_1 t_3$ & $t_2 < -\psi_2 t_3$;
(b) $t_2 \geq -\psi_2 t_3$ & $(t_1-\tau t_2)/\sqrt{t_3^2-t_2^2} < (-\psi_1+\tau\psi_2)/\sqrt{1-\psi_2^2}$;
(c) $t_1 \geq -\psi_1 t_3$ & $(t_2-\tau t_1)/\sqrt{t_3^2-t_1^2} < (-\psi_2+\tau\psi_1)/\sqrt{1-\psi_1^2}$.

Theorem 4. *The Two Hyperplane Test (THT) for the region* $\mathcal{R}(\mathbf{q},r;\mathbf{n}_1,c_1;\mathbf{n}_2,c_2)$ *is:*

$$T_{\mathcal{R}}(\mathbf{b}_i) = \begin{cases} 1, & \text{if } (a')\\ 0, & \text{otherwise;} \end{cases} \tag{39}$$

where condition (a′) is

$$V_l(\mathbf{n}_1^T\mathbf{b}_i, \mathbf{n}_2^T\mathbf{b}_i, \|\mathbf{b}_i\|_2) < \mathbf{q}^T\mathbf{b}_i < V_u(\mathbf{n}_1^T\mathbf{b}_i, \mathbf{n}_2^T\mathbf{b}_i, \|\mathbf{b}_i\|_2);$$

with $V_u(t_1,t_2,t_3) = 1 - M_2(t_1,t_2,t_3)$ *and for the lasso,* $V_l(t_1,t_2,t_3) = -V_u(-t_1,-t_2,t_3)$*, and for the nonnegative lasso,* $V_l(t_1,t_2,t_3) = -\infty$.

Theorem 4 indicates that THT uses only the $3p$ correlations $\{\mathbf{q}^T\mathbf{b}_i, \mathbf{n}_1^T\mathbf{b}_i, \mathbf{n}_2^T\mathbf{b}_i\}_{i=1}^p$. So the test has time complexity $O(pn)$.

4.5.1 Parameter selection

Assume the sphere $S(\mathbf{q},r)$ has been selected. The inequality constraints in (3) provide the natural half space bounds $\hat{\boldsymbol{\theta}} \in H(\mathbf{b})$, $\mathbf{b} \in \mathfrak{B}$. $H(\mathbf{b})$ can be equivalently specified as $\{\boldsymbol{\theta}\colon \mathbf{n}^T\boldsymbol{\theta} \leq c\}$ with $\mathbf{n} = \mathbf{b}/\|\mathbf{b}\|_2$ and $c = 1/\|\mathbf{b}\|_2$ and the resultant dome $H(\mathbf{b}) \cap S(\mathbf{q},r)$ has parameters given by (24), (25) and (26).

We seek two such half spaces. We can select the first by minimizing its dome radius r_d. By (26), this requires maximizing ψ_d:

$$\mathbf{b}^{(1)} = \underset{\mathbf{b}\in\mathfrak{B}}{\arg\max}\ \frac{\mathbf{b}^T\mathbf{q}-1}{\|\mathbf{b}\|_2}. \tag{40}$$

When all features have equal norm, we can simply maximize $\mathbf{b}^T\mathbf{q}$ over $\mathbf{b} \in \mathfrak{B}$.

Suppose we have selected the first feature $\mathbf{b}^{(1)}$ using (40). This yields a dome with dome center $\mathbf{q}^{(1)} = \mathbf{q}_d$ and dome radius $r^{(1)} = r_d$. Assume that $\psi_d \geq 0$. Then by Lemma 4, the smallest sphere containing the dome has center $\mathbf{q}^{(1)}$ and radius $r^{(1)}$. To select the second feature, we can focus on the sphere $S(\mathbf{q}^{(1)}, r^{(1)})$ and repeat the above construction:

$$\mathbf{b}^{(2)} = \underset{\mathbf{b}\in\mathfrak{B}/\mathbf{b}^{(1)}}{\arg\max}\ \frac{\mathbf{b}^T\mathbf{q}^{(1)}-1}{\|\mathbf{b}\|_2}. \tag{41}$$

When all features have equal norm, we can simply maximize $\mathbf{b}^T\mathbf{q}^{(1)}$ over $\mathbf{b} \in \mathfrak{B}/\mathbf{b}^{(1)}$. We call this parameter selection method *Dictionary-based THT* (D-THT).

Alternatively, if we have solved the instance $(\mathbf{y}_0, \lambda_0)$ yielding primal and dual solutions $\hat{\mathbf{w}}_0$ and $\hat{\boldsymbol{\theta}}_0$ (see (4)), then $\hat{\boldsymbol{\theta}}_0$ must satisfy the inequalities (7). Using some algebra and (4), these inequalities can be written as: $(B\hat{\mathbf{w}}_0)^T\boldsymbol{\theta} \leq (B\hat{\mathbf{w}}_0)^T\hat{\boldsymbol{\theta}}_0$. Since $0 \in \mathcal{F}$, the right hand side is nonnegative. Hence the inequality bounds $\mathcal{F}$ in the half space $\mathbf{n}_1^T\boldsymbol{\theta} \leq c_1$ with

$$\mathbf{n}_1 = B\hat{\mathbf{w}}_0/\|B\hat{\mathbf{w}}_0\|_2, \qquad c_1 = \mathbf{n}_1^T\hat{\boldsymbol{\theta}}_0. \tag{42}$$

One can then select $\mathbf{n}_2$ and c_2 using (41).

We will return to the THT tests in §7 where we will compare the performance of the tests with $m = 0, 1, 2$ and examine the trade-off between rejection rate and computational efficiency that increasing m imposes.

4.5.2 Connections with the Literature

The form of the Two Hyperplane Test was first presented (without proof) in [36], for unit norm features and target vector. The form given here (with proofs) is a generalization of that result. The general formulation allows the use of any sphere and hyperplane constraints bounding $\hat{\boldsymbol{\theta}}$ and includes the feature constraint used in [24] as a special case.

4.6 Composite Tests

The construction described in §4.4 gives rise to a finite sequence of spheres and domes:

$$S_1 \supset D_1 \subset S_2 \supset \cdots \supset S_{k-1} \supset D_{k-1} \subset S_k$$

Each sphere and dome has an associated test. But since D_j is contained in S_j and S_{j+1}, each dome test is stronger than the tests for the spheres that precede and succeed it. But S_{j+1} is not contained in S_j and D_{j+1} is not contained in D_j. So we can't claim that the last dome D_{k-1} leads to the strongest test. Moreover, a test based on the region $\cap_{j=1}^{k-1} D_j$ is usually too complex to compute.

An alternative is to implement a composite test that rejects $\mathbf{b}_i$ if it is rejected by any of the tests $\{T_{D_j}\}_{j=1}^{k-1}$. For the nonnegative lasso, T_{D_j} takes the form $\mu_j(\mathbf{b}_i) < 1$, with $\mu_j(\mathbf{b}_i) = \mathbf{q}_j^T\mathbf{b}_i + M_1(\mathbf{n}_j^T\mathbf{b}_i, \|\mathbf{b}_i\|_2)$ and M_1 given by (27). So the composite test rejects $\mathbf{b}_i$ if

$$\min_{j=1:k}\{\ \mathbf{q}_j\mathbf{b}_i + M_1(\mathbf{n}_j^T\mathbf{b}_i, \|\mathbf{b}_i\|_2)\ \} < 1. \tag{43}$$

Similarly, for the lasso problem the composite test rejects $\mathbf{b}_i$ if

$$\min_{j=1:k}\{\ \max\{\mathbf{q}_j^T\mathbf{b}_i + M_1(\mathbf{n}_j^T\mathbf{b}_i, \|\mathbf{b}_i\|_2), \\ -\mathbf{q}_j^T\mathbf{b}_i - M_1(-\mathbf{n}_j^T\mathbf{b}_i, \|\mathbf{b}_i\|_2\}\ \} < 1. \tag{44}$$

Reflecting the dome construction method, we call the tests (43) and (44) *iteratively refined dome tests* (IRDT). These tests can be implemented in several ways and

extra domes arising in the course of the construction can also be included. This is illustrated in §6. The major cost of the tests is calculating the inner products $\mathbf{q}_j^T\mathbf{b}_i$ and $\mathbf{n}_j^T\mathbf{b}_i$ for each feature $\mathbf{b}_i$ to be tested. Because of the iterative construction, this can be done by computing $\mathbf{q}_1^T\mathbf{b}_i, \mathbf{n}_1^T\mathbf{b}_i, \ldots, \mathbf{n}_{k-1}^T\mathbf{b}_i$ (see (34), (35), (37)). So to execute all of the tests $D_1, \ldots, D_k$, only k inner products are used per feature tested. This is $O(nk)$ time complexity per feature tested where n is the feature dimension. So the marginal cost of increasing k by 1 is the cost of computing one additional inner product per feature tested.

A composite test is mathematically equivalent to test disjunction, $(T_1 \vee T_2)(\mathbf{b}_i) = T_1(\mathbf{b}_i) \vee T_2(\mathbf{b}_i)$. The following lemma indicates that a disjunction of region tests is weaker than the test based on the intersection of the regions. It is easy to construct examples that demonstrate this. Consider two spheres of equal radius with a small intersection. Both spheres can intersect a half space while the intersection does not.

Lemma 6. *For compact sets $\mathcal{R}_1$, $\mathcal{R}_2$: $T_{\mathcal{R}_1} \vee T_{\mathcal{R}_2} \preceq T_{\mathcal{R}_1 \cap \mathcal{R}_2}$.*

Lemma 6 indicates that a disjunction of tests is trading rejection performance for simplicity and ease of implementation. Despite the above limitation, the IRDT test is very competitive with Dictionary-based THT on the data sets used in our numerical studies.

4.6.1 Connections with the Literature

The sphere test ST3 in [23] is based on a refined spherical bound. In [23] it is assumed that $\mathbf{y}$ and all features have unit norm. ST3 is then constructed starting with the default spherical bound $S(\mathbf{q}_1, r_1)$ with $\mathbf{q}_1 = \mathbf{y}/\lambda$ and $r_1 = (1/\lambda - 1/\lambda_{\max})$. The greedy strategy selects the feature $\mathbf{b} = \mathbf{b}_{\max}$. Then the dual solution $\hat{\boldsymbol{\theta}}$ lies in the default dome formed by the intersection of the spherical ball $S(\mathbf{q}_1, r_1)$ and the half space $H(\mathbf{b}_{\max})$. This intersection is indicated by the green dome region $\mathcal{G}$ in Fig. 1(c). The smallest spherical ball bounding $\mathcal{G}$ (dashed magenta circle in Fig. 1(c)) is obtained by substituting the values of $\mathbf{q}_1$, r_1 and $\mathbf{b}_{\max}$ into (34), (35) and (36). This yields $\psi_2 = \lambda_{\max}$, $\mathbf{q}_2 = \mathbf{y}/\lambda - (\lambda_{\max}/\lambda - 1)\,\mathbf{b}_{\max}$ and $r_2 = \sqrt{1/\lambda_{\max}^2 - 1}\,(\lambda_{\max}/\lambda - 1)$. These parameters are derived in [23] using a distinct approach.

A two term disjunction test is used in [26]. This test is implemented sequentially. The first test is applied and then the second is applied to the remaining features. Any disjunction test can be implemented sequentially in this fashion. The key innovation in [26] is that each test is based on an ellipsoidal bound on $\hat{\boldsymbol{\theta}}$. The first ellipsoid is the minimum volume ellipsoid containing the default dome (30). The second ellipsoid is constructed in a greedy fashion by selecting a feature so that the best ellipsoidal bound of the intersection of its half space and the first ellipsoid has minimum volume. The first step is in the spirit of ST3 except using an ellipsoidal bound. The second step is bound refinement based on ellipsoids rather than spheres. An ellipsoidal bound is tighter than the spherical bounds used in this section. However, its description requires a center $\mathbf{q} \in \mathbb{R}^n$ and a matrix $P \in \mathbb{R}^{n \times n}$ to encode its shape and orientation. When n is large this could be an impediment. In contrast, a sphere only requires specification of a center $\mathbf{q} \in \mathbb{R}^n$ and a scalar radius r.

5 Sequential Screening

The screening tests discussed so far, screen the dictionary once and then solve the reduced lasso problem that results. We hence refer to these tests as "one-shot" screening tests. These tests can perform well for moderate to large values of $\lambda/\lambda_{\max}$ but often fail to provide equivalent rejection performance for smaller values of $\lambda/\lambda_{\max}$. This is primarily due to the challenge of obtaining a tight region bound on $\hat{\boldsymbol{\theta}}$ when $\lambda/\lambda_{\max}$ is small.

Alternative screening methods can help with this problem. For example, [34] examined the idea of screening and solving (1) for a sequence of instances $\{(\mathbf{y}, \lambda_k)\}_{k=1}^N$ (Recursive-SAFE). At step k the previously solved instance $(\mathbf{y}, \lambda_{k-1})$ defines a bound on the dual solution of the instance $(\mathbf{y}, \lambda_k)$. Hence the previous solution can help screen the next instance in the sequence. A similar idea is proposed by [20] in the form of the Strong Sequential Rule. This is used to solve the lasso problem "over a grid of λ values". In [35], the SAFE test for the lasso is upgraded to use a specific dome test. Wang et al. [37], examined the design of the sequence $\{y_k\}$ on the performance of sequential screening, and showed that a geometrically spaced sequences outperform both uniform sampling and one-shot screening tests in both rejection power and computation time.

In a similar spirit, [27] proposed running a homotopy algorithm to find a solution at the K-th breakpoint on the regularization path of $\hat{\mathbf{w}}(\lambda)$. This effectively solves a sequence of lasso problems (via homotopy) to obtain a solution $\hat{\mathbf{w}}_K$ at $\lambda_K > \lambda_t$. The dual solution $\hat{\boldsymbol{\theta}}_K$ is then used to screen the instance $(\mathbf{y}, \lambda_t)$. This has potential advantages, but relinquishes control of the values λ_k used to the breakpoints in the homotopy algorithm. In the worst case the regularization path can have $O(3^p)$ breakpoints [38]. As a variant on homotopy, Sequential Lasso [39] solves a sequence of partially ℓ_1 penalized least squares problems where features with non-zero weights in earlier steps are not penalized in subsequent steps.

With the exception of homotopy, all of the above sequential schemes use a fixed open loop design for N and the sequence $\{\lambda_k\}_{k=1}^N$. For example, first fix $N \geq 2$, then select $\lambda_1 < \lambda_{\max}$, $\lambda_N = \lambda_t$, and let the intermediate values be selected via geometric spacing: $\lambda_k = \alpha\lambda_{k-1}$ with $\alpha = (\lambda_t/\lambda_1)^{1/(N-1)}$. To solve

instance $(\mathbf{y}, \lambda_t)$, we first screen and solve the instance $(\mathbf{y}, \lambda_1)$. Then sequentially for $k = 2, \ldots, N$, we screen instance $(\mathbf{y}, \lambda_k)$ with the help of the known solution of the previous instance, and then solve the reduced problem. This continues until the solution for $(\mathbf{y}, \lambda_t)$ is obtained. Usually, we are only interested in the solution of the instance $(\mathbf{y}, \lambda_t)$ – the other instances are merely waypoints in the computation.

The solution of the previous instance helps screen the next instance as follows. First form the basic bounding sphere (21) for $\hat{\boldsymbol{\theta}}_k$ with center $\mathbf{y}/\lambda_k$ and radius $\|\mathbf{y}/\lambda_k - \hat{\boldsymbol{\theta}}_{k-1}\|_2$. This uses $\hat{\boldsymbol{\theta}}_{k-1}$ as a dual feasible point. Then we use that fact that $\hat{\boldsymbol{\theta}}_{k-1}$ is the projection of $\mathbf{y}/\lambda_{k-1}$ onto $\mathcal{F}$ to form the halfspace (31) with $\mathbf{n}_{k-1} = (\mathbf{y}/\lambda_{k-1} - \hat{\boldsymbol{\theta}}_{k-1})/\|\mathbf{y}/\lambda_{k-1} - \hat{\boldsymbol{\theta}}_{k-1}\|_2$, and $c_{k-1} = \mathbf{n}_{k-1}^T \hat{\boldsymbol{\theta}}_{k-1}$. This sphere and halfspace yield the bounding dome derived in §4.3.2:

$$D_k = D(\mathbf{y}/\lambda_k, r_k; \mathbf{n}_{k-1}, c_{k-1}) \tag{45}$$

This dome, illustrated in Fig. 4, encapsulates information about $\hat{\boldsymbol{\theta}}_k$ provided by the dual solution of the previous instance.

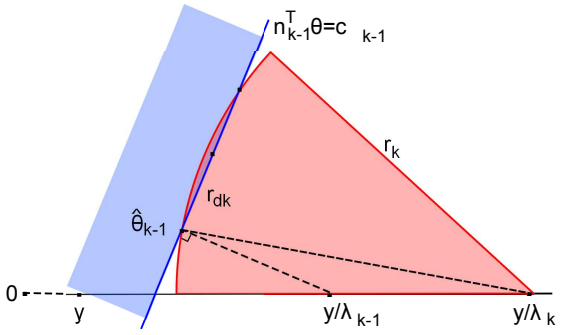

Fig. 4. An illustration of the dome (45) formed at step k.

In contrast to an open loop design, one can use feedback to adaptively select N and the sequence $\{\lambda_k\}_{k=1}^N$ as the computation proceeds [40]. This allows the value of N and the sequence $\{\lambda_k\}_{k=1}^N$ to be adapted to each particular instance. For some instances, a small value of N is used, for others, a larger value is used. This Data-Adaptive Sequential Screening (DASS) scheme has the potential for greater robustness and efficiency over a fixed open loop design. One way to see why feedback helps is to examine the diameter of D_k in (45). For this we use the following proposition from [40].

Proposition 1. *Let D_k be the dome* (45) *and $\delta_k = diam(D_k)$. Then*

$$\delta_k = 2\left(\frac{1}{\lambda_k} - \frac{1}{\lambda_{k-1}}\right)\sqrt{\mathbf{y}^T(I - \mathbf{n}_{k-1}\mathbf{n}_{k-1}^T)\mathbf{y}}. \tag{46}$$

Using Proposition 1, it can be shown that the feedback selection rule

$$\frac{1}{\lambda_k} = \frac{1}{\lambda_{k-1}} + \frac{1/2R}{\sqrt{\mathbf{y}^T(I - \mathbf{n}_{k-1}\mathbf{n}_{k-1}^T)\mathbf{y}}}, \tag{47}$$

where $R > 0$ is a selectable parameter, ensures that $\text{diam}(D_k) \leq R$ for all $k > 1$. This allows direct control of how tightly the dome (45) bounds $\hat{\boldsymbol{\theta}}_k$, $k = 1, \ldots, N$. Note that in this scheme N is not predetermined. Instead the stopping time is decided by the feedback scheme. However, (47) ensures that $N \leq 1 + \log(1/\lambda_t)/\log(1 + C/2R)$, where C is an bound on the dual regularization path [40]. We employ this feedback method to demonstrate the effectiveness of sequential screening in §7.

6 Algorithms

Each of the screening tests previously described requires the inputs B, $\mathbf{y}$, and λ and returns $v_1, \ldots, v_p$ where v_i is a logical value indicating if $\mathbf{b}_i$ is rejected. The algorithms can be implemented in an online fashion with very few features stored in memory at once. The critical computation is calculating inner products of the form $\mathbf{y}^T\mathbf{b}_i$ and $\mathbf{n}^T\mathbf{b}_i$. If the features are sparse or if specialized hardware is available for such computations, then running times are further reduced.

A basic implementation of THT is shown in Algorithm 1. If the dictionary is unnormalized, the feature norms can be precomputed and passed as an input to the algorithm. If it is normalized, we recommend simplifying the algorithm by setting $\beta_i = 1$ and removing unnecessary floating point operations (see §4.5). The algorithm accepts two additional optional inputs: either a dual solution $(\boldsymbol{\theta}_F, \lambda_2)$ or only a feasible point $\boldsymbol{\theta}_F$. The dual solution is useful for the application of THT in sequential screening. It is used to select the first half space used by THT (lines 7,8). If only a feasible point is provided, it is used to select the sphere radius. Otherwise, the default point $\mathbf{y}/\lambda_{\max}$ is used. The remaining half spaces are selected using dictionary-based selection (40), (41) (§4.5). The output values v_i are determined for each $\mathbf{b}_i$ by evaluation of the THT test in Theorem 4.

A basic implementation of IRDT is shown in Algorithm 2. To keep the notation simple and the algorithm understandable, all features and $\mathbf{y}$ are assumed to have unit norm, but this is not required (see §4.6). IRDT uses at most s iterations with the value of s supplied by the user (we recommend $s \leq 5$). The algorithm passes through the dictionary at most $s+1$ times with the main loop executed at most s times. The break at line 15 terminates this loop early if suitable domes can't be found. The algorithm accepts a feasible point $\boldsymbol{\theta}_F$ for the dual problem as an optional input and can be adapted to accept a known dual solution. The value v_i for each $\mathbf{b}_i$ is determined by a disjunction of a set of dome tests each based on the dome test in Theorem 3. These disjunctions are

Algorithm 1 Two Hyperplane Test (THT)

Input: Required: $\{\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_p\}$, $\mathbf{y}$, λ.
Optional: $\{\beta_1, \ldots, \beta_p\}$ with $\beta_i = \|\mathbf{b}_i\|_2$;
Optional: $\boldsymbol{\theta}_F \in \mathcal{F}$;
Optional: $(\lambda_2, \boldsymbol{\theta}_F)$ a dual solution.
Output: $v_1, v_2, \ldots, v_p$ (if $v_i = 1$, $\mathbf{b}_i$ is rejected).

if $\{\beta_i, i = 1, \ldots, p\}$ is not provided **then**
  $\beta_i \leftarrow \|\mathbf{b}_i\|_2, i = 1, \ldots, p$
end if
$\mathbf{q} \leftarrow \mathbf{y}/\lambda$. (sphere center)
$\rho_i \leftarrow \mathbf{q}^T\mathbf{b}_i, 1 \le i \le p$.
if a dual solution is provided **then**
  $\mathbf{n}_1 \leftarrow (\mathbf{y}/\lambda_2 - \boldsymbol{\theta}_F)/\|\mathbf{y}/\lambda_2 - \boldsymbol{\theta}_F\|_2$.
  $c_1 \leftarrow \mathbf{n}_1^T\boldsymbol{\theta}_F$.
else
  if $\boldsymbol{\theta}_F$ is not provided **then**
    $\lambda_{\max} \leftarrow \lambda \max_i\{f(\rho_i)\}$.
    $\boldsymbol{\theta}_F \leftarrow \mathbf{y}/\lambda_{\max}$.
  end if
  $i_* \leftarrow \arg\max_i\{(f(\rho_i) - 1)/\beta_i\}$.
  $\mathbf{n}_1 \leftarrow \mathbf{b}_{i_*}/\beta_{i_*}$.
  $c_1 \leftarrow 1/\beta_{i_*}$.
end if
$r \leftarrow \|\boldsymbol{\theta}_F - \mathbf{y}/\lambda\|_2$. (sphere radius)
$a \leftarrow \mathbf{n}_1^T\mathbf{q} - c_1$.
$\sigma_i \leftarrow \mathbf{n}_1^T\mathbf{b}_i, 1 \le i \le p$.
$t_i \leftarrow \rho_i - a\sigma_i, 1 \le i \le p$.
if a dual solution is provided **then**
  $j_* \leftarrow \arg\max_i\{(f(t_i) - 1)/\beta_i\}$.
else
  $j_* \leftarrow \arg\max_{i \ne i_*}\{(f(t_i) - 1)/\beta_i\}$.
end if
$\mathbf{n}_2 \leftarrow \mathbf{b}_{j_*}/\beta_{j_*}$.
$c_2 \leftarrow 1/\beta_{j_*}$.
$\tau_i \leftarrow \mathbf{n}_2^T\mathbf{b}_i, 1 \le i \le p$.
$v_i \leftarrow [\![V_l(\sigma_i, \tau_i, \beta_i) < \rho_i < V_u(\sigma_i, \tau_i, \beta_i)]\!]$

Fig. 5. Algorithm for THT. The functions V_u and V_l are from Theorem 4. **Other Notation:** For the lasso, $f(z) = |z|$ and $g(z) = \text{sign}(z)$ and for the nonnegative lasso, $f(z) = g(z) = z$. For a logical condition $c(\cdot)$, $[\![c(z)]\!]$ evaluates to *true* if z satisfies condition c and *false* otherwise.

Algorithm 2 Iteratively Refined DT

Input: Required: $\{\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_p\}$, $\mathbf{y}$, λ, s
For simplicity, assume $\|\mathbf{y}\|_2 = \|\mathbf{b}_i\|_2 = 1$.
Optional: $\boldsymbol{\theta}_F \in \mathcal{F}$.
Output: $v_1, v_2, \ldots, v_p$ (if $v_i = 1$, $\mathbf{b}_i$ is rejected).

$\mathbf{q}_1 \leftarrow \mathbf{y}/\lambda$.
$\rho_{i,1} \leftarrow \mathbf{q}_1^T\mathbf{b}_i, 1 \le i \le p$.
if $\boldsymbol{\theta}_F$ is not provided **then**
  $\boldsymbol{\theta}_F \leftarrow \mathbf{y}/(\lambda \max_i f(\rho_{i,1}))$.
end if
$r_1 \leftarrow \|\boldsymbol{\theta}_F - \mathbf{y}/\lambda\|_2$.
$v_i \leftarrow [\![f(\rho_{i,1}) < 1 - r_1]\!], 1 \le i \le p$.
$\sigma_i \leftarrow$ *false*, $1 \le i \le p$
for $j_1 = 1, 2, \ldots, s$ **do**
  $h \leftarrow \arg\max_{v_i = \textit{false}, \sigma_i = \textit{false}} f(\rho_{i,j_1})$.
  $\mathbf{b} \leftarrow g(\rho_{h,j_1})\mathbf{b}_h$.
  $t_i \leftarrow \mathbf{b}^T\mathbf{b}_i, 1 \le i \le p$.
  $\psi \leftarrow (f(\rho_{h,j_1}) - 1)/r_{j_1}$.
  if $\psi \le 0$ **then**
    BREAK.
  end if
  if $j_1 < s$ **then**
    $\mathbf{q}_{j_1+1} \leftarrow \mathbf{q}_{j_1} - \psi r_{j_1}\mathbf{b}$.
    $\rho_{i,j_1+1} \leftarrow \rho_{i,j_1} - \psi r_{j_1} t_i, 1 \le i \le p$.
    $r_{j_1+1} \leftarrow r_{j_1}\sqrt{1 - \psi^2}$.
  end if
  for $j_2 = j_1, j_1 - 1, \ldots, 1$ **do**
    if $j_2 < j_1$ **then**
      $\psi \leftarrow (\mathbf{q}_{j_2}^T\mathbf{b} - 1)/r_{j_2}$.
    end if
    $r \leftarrow r_{j_2}$.
    for $i \in \{i : v_i =$ *false*$\}$ **do**
      $v_i \leftarrow [\![V_l(t_i) < \rho_{i,j_2} < V_u(t_i)]\!]$
    end for
  end for
  $\sigma_h \leftarrow$ *true*.
end for

Fig. 6. Algorithm for IRDT. The functions V_u and V_l are from Theorem 3. See the caption of Algorithm 1 for the definition of other notation.

computed sequentially with subsequent tests applied only to currently surviving features.

Data-Adaptive Sequential Screening solves N lasso instances $\{(\mathbf{y}, \lambda_k)\}_{k=1}^N$ for a sequence of descending values λ_k where $\lambda_{\max} > \lambda_1$ and $\lambda_N = \lambda_t$ is the regularization parameter value for the instance to be solved. The user must specify a radius $R > 0$. At each step, the algorithm uses a strong "one-shot" screening test, for example THT, provided with a solution of the previous instance, followed by an external lasso solver to solve the screened current instance. The algorithm sets $\lambda_1 = 0.95\lambda_{\max}$ and thereafter uses the feedback rule (47) to select λ_k until $\lambda_k \le \lambda_t$. It then sets $N = k$, $\lambda_N = \lambda_t$ and screens and solves the final problem. See [40] for additional details on this algorithm.

7 Numerical Examples

We now empirically examine the performance of the screening algorithms presented using the datasets summarized in Table 1, and discussed in detail below.
(1) RAND: We generate lasso problems with $p = 10,000$ features of dimension $n = 28$ by randomly generating $10,001$ 28-dimensional vectors $\mathbf{y}, \mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_{10,000}$ using the MATLAB $rand()$ function. These vectors are then scaled to unit norm.
(2) MNIST500: The MNIST data set [41], [42], consists of $70,000$ 28×28 images of hand-written digits ($60,000$ in the training set, $10,000$ in the testing set). We form a dictionary by randomly sampling 500 images for each digit from the training set, and a target vector by randomly sampling an image from the testing set. Each image is vectorized and scaled to unit norm.
(3) YALEBXF: This data set contains the frontal face images (192×168) of the 38 subjects in the extended

Algorithm 3 Data-Adaptive Sequential Screening

Input: $\{\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_p\}$, $\mathbf{y}$, λ_t, $R > 0$, a lasso solver $\mathbb{S}$. For simplicity, assume $\|\mathbf{y}\|_2 = \|\mathbf{b}_i\|_2 = 1$.
Output: $\hat{\mathbf{w}}_t$

1: $\lambda_{\max} \leftarrow \max_i |\mathbf{y}^T \mathbf{b}_i|$.
2: $k \leftarrow 1$, $\lambda_1 \leftarrow 0.95\lambda_{\max}$.
3: call THT with $\{\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_p\}$, $\mathbf{y}$, λ_1 only.
4: call $\mathbb{S}$ to solve the lasso problem $(\mathbf{y}, \lambda_1)$ using the non-rejected features, get $\hat{\mathbf{w}}_1$.
5: $\hat{\boldsymbol{\theta}}_1 \leftarrow (\mathbf{y} - [\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_p]\hat{\mathbf{w}}_1)/\lambda_1$.
6: **while** $\lambda_k > \lambda_t$ **do**
7: $k \leftarrow k + 1$.
8: $\mathbf{n}_{k-1} \leftarrow \frac{\mathbf{y}/\lambda_{k-1} - \hat{\boldsymbol{\theta}}_{k-1}}{\|\mathbf{y}/\lambda_{k-1} - \hat{\boldsymbol{\theta}}_{k-1}\|_2}$.
9: $\frac{1}{\lambda_k} \leftarrow \frac{1}{\lambda_{k-1}} + \frac{1/2R}{\sqrt{\mathbf{y}^T(I - \mathbf{n}_{k-1}\mathbf{n}_{k-1}^T)\mathbf{y}}}$.
10: **if** $\lambda_k < \lambda_t$ **then**
11: $\lambda_k \leftarrow \lambda_t$.
12: **end if**
13: call THT with $\{\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_p\}$, $\mathbf{y}$, λ_k, and a dual solution $(\lambda_{k-1}, \hat{\boldsymbol{\theta}}_{k-1})$.
14: call $\mathbb{S}$ to solve the lasso problem $(\mathbf{y}, \lambda_k)$ using the non-rejected features, get $\hat{\mathbf{w}}_k$.
15: $\hat{\boldsymbol{\theta}}_k \leftarrow (\mathbf{y} - [\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_p]\hat{\mathbf{w}}_j)/\lambda_k$.
16: **end while**
17: $\hat{\mathbf{w}}_t \leftarrow \hat{\mathbf{w}}_k$.

Fig. 7. Algorithm for Data-Adaptive Sequential Screening.

Yale B face data set [43], [44]. We randomly select $p = 2,000$ of the $2,414$ face images as the dictionary, and randomly pick $\mathbf{y}$ from the remaining 414 images. Each image is vectorized and scaled to unit norm.

(4) NYT: This is a bag-of-words dataset from the UCI Machine Learning Repository [45], in which 300,000 New York Times articles are represented as vectors with respect to a vocabulary of 102,660 words. The i-th entry in vector j gives the number of occurrences of word i in document j. Documents with low word counts are removed, leaving 299,752 documents. We select the first 299,000 of these as the dictionary and 6 documents from the remaining 752 as target vectors.

All experiments solve the standard lasso problem (1) using the Feature-sign lasso solver [46]. The grafting lasso solver [47] was also tested and gave similar qualitative performance. We use two performance metrics: the percentage of features rejected and the speedup in solving the lasso problem. Recall, speedup is the time to solve the lasso problem divided by sum of the time to screen *and* the time to solve the reduced lasso problem. Timing and speedup results depend on the solver used. The regularization parameter λ is set using the scaling invariant ratio $\lambda/\lambda_{\max}$ where $\lambda_{\max} = \max_i |\mathbf{y}^T \mathbf{b}_i|$. So $\lambda/\lambda_{\max} \in [0, 1]$ with larger values yielding sparser solutions. Except for NYT, we randomly select 64 lasso problems and report the average result and the standard error.

Data Set	# features p	Dim. n	Av. $\lambda_{\max}$ (stnd. err.)
RAND	10,000	28	0.919 (0.002)
MNIST500	5,000	784	0.865 (0.005)
YALEBXF	2,000	32,256	0.963(0.008)
NYT	299,000	102,660	0.714 (-)

TABLE 1
$\lambda_{\max}$ is averaged over the lasso instances solved.

7.1 The performance of one-shot screening

We first benchmark the performance of the one-shot tests: ST (§4.2), DT (§4.3), and D-THT (§4.5). To begin, we use the default spherical bound (22). This conservative choice gives a lower bound for the performance of the one-shot screening methods on each dataset. The default dome test combines this sphere with the feature $\mathbf{b}_{\max}$, while dictionary-based THT combines it with two features using the selection scheme detailed in (40), (41). We also show results using a second "oracle" bounding sphere with center $\mathbf{y}/\lambda$ and radius $r = \|\mathbf{y}/\lambda - \hat{\boldsymbol{\theta}}\|_2$. This bound is not practical, but provides a useful upper bound on performance over bounding spheres centered at $\mathbf{y}/\lambda$.

The performance of the one-shot screening methods on the test datasets is shown in Fig. 8. Here are the salient points: (a) While the default one-shot tests perform well for high values of $\lambda/\lambda_{\max}$, this performance quickly degrades as $\lambda/\lambda_{\max}$ decreases. At values of $\lambda/\lambda_{\max}$ around 0.2 and lower, the tests are not effective. (b) On the other hand, the upper bounds indicate potential for improvement if a better spherical bound can be found. Indeed, the significant gap between the lower and upper performance bounds indicates that it is worth investing computation to improve the default spherical bound. (c) Among the tested methods, D-THT exhibits the best performance except at very high values of $\lambda/\lambda_{\max}$. On RAND, for example, using $\lambda/\lambda_{\max} = 0.5$ and the default spherical bound, D-THT yields a 400% rejection improvement over DT. The concurrent speedup for D-THT is about 5X while for DT is less than 2X. These effects are also seen for MNIST500 and YALEBXF.

The rejection and speedup of IRDT (not plotted) and D-THT were very similar on the test datasets with IRDT terminating after 3 or 4 iterations at the break in line 14-16 in Algorithm 2. Only for YALEXBF does D-THT pull slightly ahead.

The results confirm our claim that current one-shot screening can perform well at moderate to high values of $\lambda/\lambda_{\max}$ but such performance does not extend to low values of $\lambda/\lambda_{\max}$.

7.2 The performance of sequential screening

To explore the effectiveness of sequential screeing, we tested the Data-Adaptive Sequential Screening (DASS) scheme (47). The performance results are shown in Fig. 9. Here are the salient points: (a) For

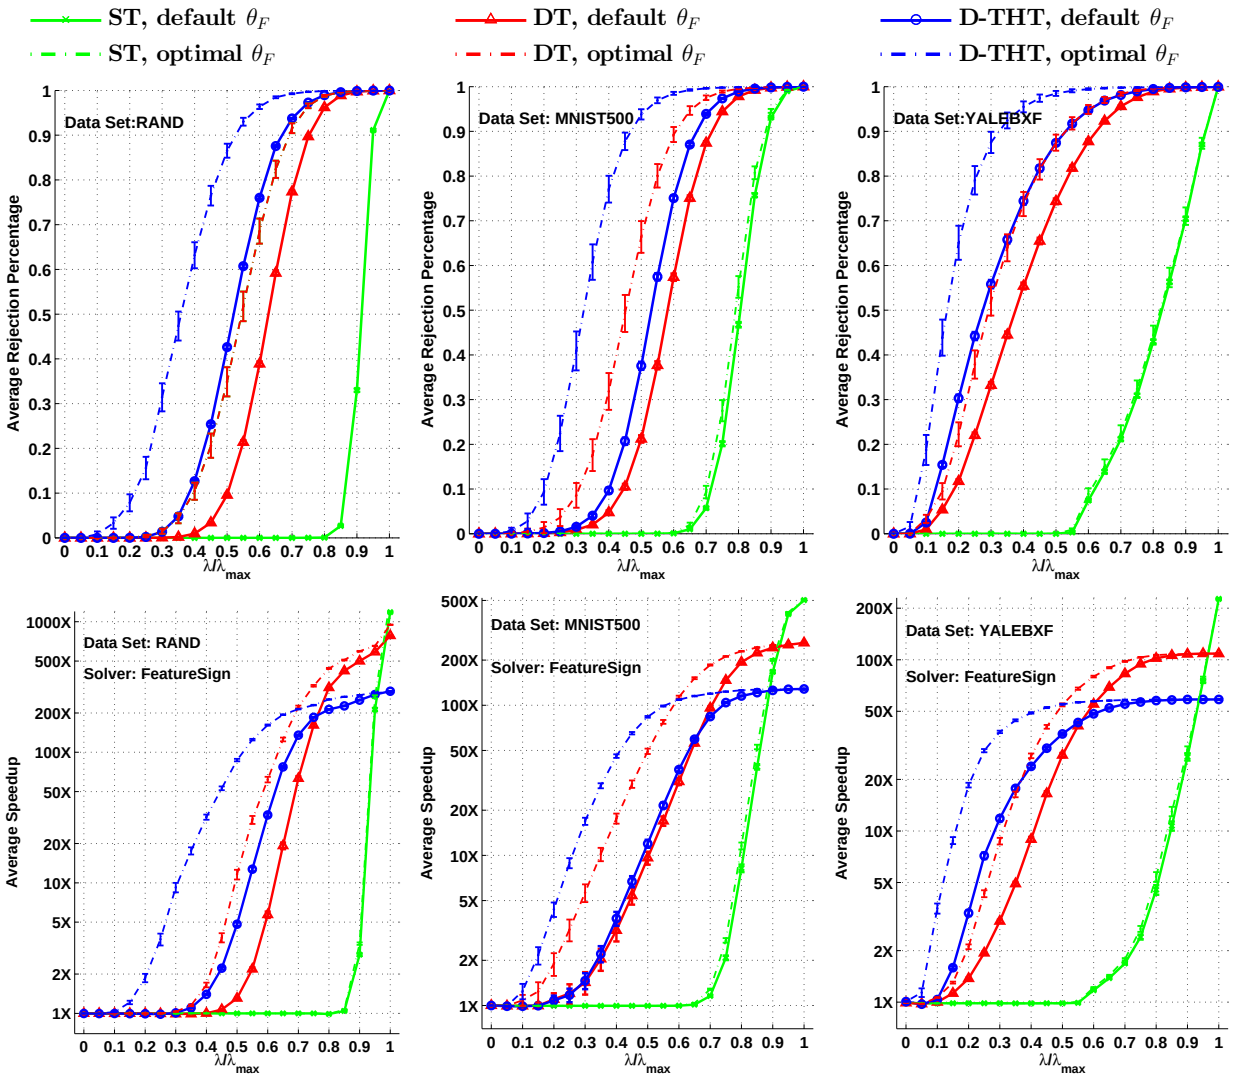

Fig. 8. Performance of ST, DT and D-THT. Top: rejection percentage; Bottom: speedup using screening and the FeatureSign solver. Solid curves lower bound and dashed curves upper bound performance for spherical bounds centered at $\mathbf{y}/\lambda$.

both MNIST500 and YALEBXF, with $R = 0.2$ the performance of DASS is robust across a variety of values of λ_t; (b) DASS yields significant improvement in rejection fraction and robust speedup performance compared with one-shot tests; (c) At values of $\lambda/\lambda_{\max}$ around 0.1 and lower, DASS is rejecting 98% of the dictionary while giving speedup greater than 1. This is successful screening at much lower values of $\lambda/\lambda_{\max}$.

7.3 Sequential screening on a large dataset

As a final experiment, we used the NYT dataset to explore how successfully DASS can screen and solve lasso problems using small values of $\lambda/\lambda_{\max}$ with high dimensional data, and a very large dictionary. We normalize each document and selected six documents from the test set to be used as target vectors. These were randomly chosen from the first 100 of the 752 held out documents subject to $0.5 < \lambda_{\max} < 0.9$. We ran screening in an "on-line" mode by loading only small amounts of the dictionary into memory at a time. We used $\lambda_t/\lambda_{\max} = 0.1$, $\lambda_1 = 0.95\lambda_{\max}$ and $R = 0.3$. The value of N is selected automatically for each instance. In one case, $N = 16$ while in another $N = 27$. In all tested cases, $N \leq 27$. As a benchmark, we also tested a geometrically spaced, open loop sequential screening algorithm using $\lambda_t/\lambda_{\max} = 0.1$, $\lambda_1 = 0.95\lambda_{\max}$ and $N = 30$.

The results for both methods are shown in Fig. 10. Due to fixed computational resources, we could not solve these lasso problems without effective screening. Hence the usual speedup metric can't be evaluated. The main time cost is sequentially reading features from disk into RAM. Here at the main points to note: (1) Under geometric spacing with fixed N, less than 10,000 of the features (3.3%) were held in memory at once; (2) For DASS, less than 1,000 of the features (0.33%) were held in memory at once – an order of magnitude improvement over fixed geometric spacing

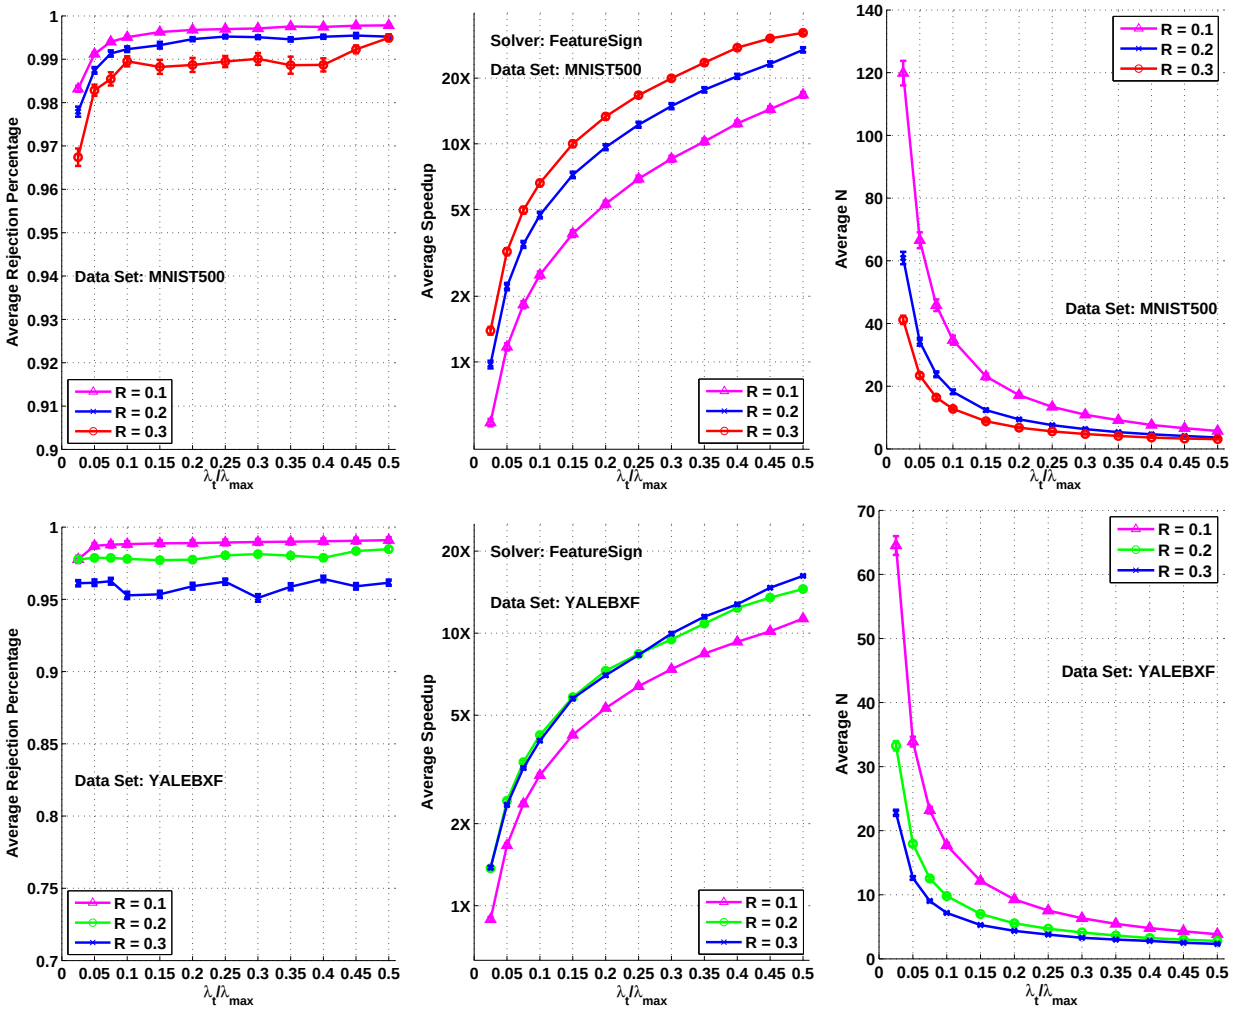

Fig. 9. Data-Adaptive Sequential Screening (DASS) applied to MNIST (top) and YALEBXF (bottom). (Left): average rejection percentage. (Middle): Speedup factor. (Right): The average value of N over the 64 trials.

(The small dip at λ_t is due to termination method); (3) On this dataset, both open loop sequential screening and DASS clearly exhibit a significant performance advantage over one-shot tests. The use of feedback by DASS to automatically select the number of steps N and the values $\{\lambda_k\}_{k=1}^{N}$, yields robust rejection performance. By tweaking N for each test vector in the open loop scheme, one could improve its average performance. But DASS handles this automatically and robustly over a wide range of values of $\lambda_t/\lambda_{\max}$.

8 Discussion and Conclusion

In our survey of key methods for screening the lasso problem, we have emphasized separating the discussion of test structure from the problem of selecting its parameters. This allowed us to see connections between many existing screening tests, and enables a clearer understanding of screening in general. Hopefully this will be advantageous to the development of new tests and parameter selection methods.

For one-shot screening tests, our numerical studies on THT strongly suggest that more complex region tests are indeed worthwhile. THT gave significant performance improvement beyond simpler tests in both rejection and speedup over important ranges of $\lambda/\lambda_{\max}$ values. The Two Hyperplane Test (THT) exhibits the best performance among one-shot tests at interesting values of $\lambda/\lambda_{\max}$ on the test datasets. The numerical studies also indicated a significant performance gap between using the default spherical bound and the best bound at the same sphere center. This indicates the value of additional computation to improve the spherical bound. This can be done by solving one or more lasso problems for the same target vector $\mathbf{y}$ but higher values of λ, or using homotopy as in [27].

Our empirical studies have also shown that sequential screening (for example, DASS) can significantly extend useful screening performance to a wider range of $\lambda/\lambda_{\max}$ (for example, as small as 0.05) than that offered by the best one-shot tests. DASS has the additional advantage that it selects both the number

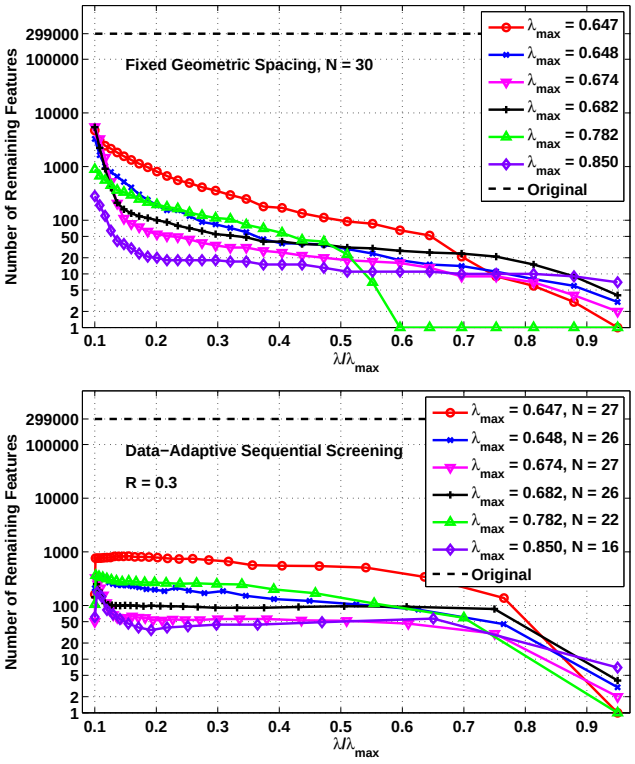
Fixed Geometric Spacing, N = 30
Number of Remaining Features
λ/λmax
λmax = 0.647
λmax = 0.648
λmax = 0.674
λmax = 0.682
λmax = 0.782
λmax = 0.850
Original
Data−Adaptive Sequential Screening
R = 0.3
λmax = 0.647, N = 27
λmax = 0.648, N = 26
λmax = 0.674, N = 27
λmax = 0.682, N = 26
λmax = 0.782, N = 22
λmax = 0.850, N = 16

Fig. 10. Sequential screening on the NYT dataset. (Top): Open loop geometric spacing with $N = 30$; (Bottom): Data-Adaptive Sequential Screening with $R = 0.3$. The graph shows the solution of six problems $(\mathbf{y}_i, \lambda_t)$, $i = 1, \ldots, 6$ with $\lambda_t = 0.1\lambda_{\max}$. The plotted points indicate the number of surviving features after THT screening of each instance $(\mathbf{y}_i, \lambda_{ik})$, $k = 1, \ldots, N$, $i = 1, \ldots, 6$.

N and the sequence $\{\lambda_k\}_{k=1}^N$ automatically.

Screening is critical when the dictionary will not fit into available memory. We have demonstrated a successful application of DASS to a very large NYT dataset, of dimension $102,660$ by $299,000$. To the best of the authors' knowledge, with constrained computational resources, screening is the only way to solve lasso problems of this size.

The concepts described in this survey should provide a firm foundation for understanding screening for related sparse representation problems. This includes screening for the elastic net (reducible to lasso problem), ℓ_1 regularized logistic regression, the graphical lasso, and the group lasso [20]. In addition, SAFE methods have been developed for the sparse support vector machine and logistic regression in [22], and the group lasso in [27]. Recently, Liu et al. [48] have proposed safe screening for generalized sparse linear models. This makes use of the variational inequality that provides a necessary and sufficient optimality condition for the dual problem. Dash et al., [49], consider screening for Boolean compressed sensing in which the objective is to select a sparse set of Boolean rules that are predictive of future outcomes. One of the screening rules developed is based on the duality arguments presented here.

Screening has also been applied in several contexts. For example, Wang et al., [50], have integrated DASS with sparse representation classification, [7], to speed up classification; Jao et al., [51], have applied screening to the problem of representing music in terms of an audio codebook (dictionary) for genre tagging; and for music genre classification, Xu et al. [52] applied screening to a large dictionary of second order scattering features. We expect to see more such applications as the size of dictionaries increase.

Appendix A
Proofs §2

The dual lasso problem (3) is obtained as follows. Setting $\mathbf{z} = \mathbf{y} - B\mathbf{w}$ in (1) gives the constrained problem: $\min_{\mathbf{z}} {}^1\!/_2\, \mathbf{z}^T\mathbf{z} + \lambda\|\mathbf{w}\|_1$, subject to $\mathbf{z} = \mathbf{y} - B\mathbf{w}$. Form the Lagrangian $\mathcal{L}(\mathbf{z}, \mathbf{w}, \boldsymbol{\mu}) = {}^1\!/_2\, \mathbf{z}^T\mathbf{z} + \lambda\|\mathbf{w}\|_1 + \boldsymbol{\mu}^T(\mathbf{y} - B\mathbf{w})$ and compute the subdifferentials with respect to $\mathbf{z}$ and $\mathbf{w}$. Using the condition that $\mathbf{0}$ must be in each subdifferential gives $\boldsymbol{\mu} = \hat{\mathbf{z}}$ and the constraints $|\boldsymbol{\mu}^T\mathbf{b}_i| \le \lambda$, $i = 1, \ldots, p$. The above equations allow the elimination of $\mathbf{z}$ and $\mathbf{w}$ from $\mathcal{L}$. This leads to the dual problem: $\max_{\boldsymbol{\mu}} {}^1\!/_2\, \|\mathbf{y}\|_2^2 - {}^1\!/_2\|\boldsymbol{\mu} - \mathbf{y}\|_2^2$, subject to $|\boldsymbol{\mu}^T\mathbf{b}_i| \le \lambda$, $i = 1, \ldots, p$. The change of variable $\boldsymbol{\theta} = \boldsymbol{\mu}/\lambda$ then gives (3). By construction, the primal and dual solutions $\hat{\mathbf{w}}$ and $\hat{\boldsymbol{\theta}}$ are related through (4).

Appendix B
Proofs §3

Theorem 1: The proof for the lasso is given above the theorem statement. For the nn-lasso, only in the definition of the active set changes. □

Corollary 1: In the proof of Theorem 1, the inclusion $\bar{S} \subseteq \bar{A}(\hat{\boldsymbol{\theta}})$, gives $i \in \bar{S}$ implies $\hat{w}_i = 0$. □

Appendix C
Proofs §4

Corollary 2: This is proved above the corollary. □

Lemma 1: Assume $\mathcal{R}_1 \subseteq \mathcal{R}_2$. If $\mathcal{R}_1 = \emptyset$, the result is clear. Hence assume $\mathcal{R}_1 \neq \emptyset$. Note that $\mu_{\mathcal{R}_1}(\mathbf{b}) = \max_{\boldsymbol{\theta} \in \mathcal{R}_1} \boldsymbol{\theta}^T\mathbf{b} \le \max_{\boldsymbol{\theta} \in \mathcal{R}_2} \boldsymbol{\theta}^T\mathbf{b} = \mu_{\mathcal{R}_2}(\mathbf{b})$. For the lasso, if $\mathbf{b}_i$ is rejected by $T_{\mathcal{R}_2}$, then $\mu_{\mathcal{R}_2}(\mathbf{b}_i) < 1$ and $\mu_{\mathcal{R}_2}(-\mathbf{b}_i) < 1$. Hence $\mu_{\mathcal{R}_1}(\mathbf{b}_i) < 1$ and $\mu_{\mathcal{R}_1}(-\mathbf{b}_i) < 1$. So $\mathbf{b}_i$ is also rejected by $T_{\mathcal{R}_1}$. Therefore $T_{\mathcal{R}_2} \preceq T_{\mathcal{R}_1}$. The proof for the nn-lasso is similar. □

Appendix D
Proofs §4.2

Lemma 2: By Cauchy-Schwarz: $\boldsymbol{\theta}^T\mathbf{b} = (\boldsymbol{\theta} - \mathbf{q})^T\mathbf{b} + \mathbf{q}^T\mathbf{b} \le \|\boldsymbol{\theta} - \mathbf{q}\|_2\|\mathbf{b}\|_2 + \mathbf{q}^T\mathbf{b}$ with equality when $\boldsymbol{\theta} - \mathbf{q}$ is aligned with $\mathbf{b}$. Then $\boldsymbol{\theta} \in S(\mathbf{q}, r)$ ensures $\boldsymbol{\theta}^T\mathbf{b} \le r\|\mathbf{b}\|_2 + \mathbf{q}^T\mathbf{b}$ with equality when $\|\boldsymbol{\theta} - \mathbf{q}\|_2 = r$. □

Theorem 2: For the nn-lasso we reject $\mathbf{b}_i$ if $\mu_S(\mathbf{b}_i) < 1$. So (19) follows from (18). For the lasso

we reject $\mathbf{b}_i$ if $\mu_S(\mathbf{b}_i) < 1$ and $\mu_S(-\mathbf{b}_i) < 1$, i.e., if $\mathbf{q}^T\mathbf{b}_i < (1 - r\|\mathbf{b}_i\|_2)$ and $\mathbf{q}^T\mathbf{b}_i > -(1 - r\|\mathbf{b}_i\|_2)$. This gives (19). Note $\max\{\mu_S(\mathbf{b}_i), \mu_S(-\mathbf{b}_i)\} < 1 \Leftrightarrow \max\{\mathbf{q}^T\mathbf{b}_i + r\|\mathbf{b}_i\|_2, -\mathbf{q}^T\mathbf{b}_i + r\|\mathbf{b}_i\|_2\} < 1 \Leftrightarrow |\mathbf{q}^T\mathbf{b}_i| < 1 - r\|\mathbf{b}_i\|_2$. Thus (19) and (20) are equivalent. □

APPENDIX E
PROOFS §4.3

Lemma 3: Solving (16) with $m = 1$ is equivalent to solving the Lagrangian problem:

$$\max_{\mu,\sigma\geq 0} \min_{\boldsymbol{\theta}} \quad \mathcal{L}(\boldsymbol{\theta}, \mu, \sigma) = -\boldsymbol{\theta}^T\mathbf{b} + \mu[(\boldsymbol{\theta} - \mathbf{q})^T(\boldsymbol{\theta} - \mathbf{q}) - r^2] + \sigma(\mathbf{n}^T\boldsymbol{\theta} - c). \tag{48}$$

Setting the derivative w.r.t. $\boldsymbol{\theta}$ equal to zero yields $\boldsymbol{\theta} = \mathbf{q} + \mathbf{b}/2\mu - \sigma\mathbf{n}/2\mu$. Substituting $\boldsymbol{\theta}$ into μ_D and (48):

$$\mu_D(\mathbf{b}) = \mathbf{b}^T\boldsymbol{\theta} = \mathbf{b}^T\mathbf{q} + \|\mathbf{b}\|_2^2/2\mu - \sigma\mathbf{b}^T\mathbf{n}/2\mu, \tag{49}$$

$$\mathcal{L}(\mu, \sigma) = -\mathbf{q}^T\mathbf{b} - \tfrac{\|\mathbf{b}\|^2}{4\mu} - \tfrac{\sigma^2}{4\mu} - \mu r^2 + \sigma r\psi + \tfrac{\sigma t}{2\mu}, \tag{50}$$

where $\psi = (\mathbf{q}^T\mathbf{n} - c)/r$ and $t = \mathbf{n}^T\mathbf{b}$. We now minimize this expression over $\mu, \sigma \geq 0$. Setting the derivatives of $\mathcal{L}$ w.r.t. μ, and σ equal to 0 yields two equations to solve for μ and σ: $\|\mathbf{b}\|_2^2 + \sigma^2 - 4\mu^2r^2 = 2\sigma t$ and $\sigma = 2\mu r\psi + t$. There are two cases: (A) If $t \geq -\psi\|\mathbf{b}\|_2$, then $\sigma = t + \psi\sqrt{\frac{\|\mathbf{b}\|_2^2 - t^2}{1-\psi^2}}$ and $\mu = \frac{1}{2r}\sqrt{\frac{\|\mathbf{b}\|_2^2 - t^2}{1-\psi^2}}$; and (B) If $t < -\psi\|\mathbf{b}\|_2$, then $\sigma = 0$ and $\mu = \|\mathbf{b}\|_2/(2r)$. Substitution of these expressions into (49) yields the result in Lemma 3. □

Theorem 3: For the nn-lasso, we reject $\mathbf{b}$ if $\mu_D(\mathbf{b}) = \mathbf{q}^T\mathbf{b} + M_1(\mathbf{n}^T\mathbf{b}, \|\mathbf{b}\|_2) < 1$, i.e., if $\mathbf{q}^T\mathbf{b} < 1 - M_1(\mathbf{n}^T\mathbf{b}, \|\mathbf{b}\|_2) = V_u(\mathbf{n}^T\mathbf{b}, \|\mathbf{b}\|_2)$. For the lasso we reject $\mathbf{b}$ if $\mathbf{q}^T\mathbf{b} + M_1(\mathbf{n}^T\mathbf{b}, \|\mathbf{b}\|_2) < 1$ and $-\mathbf{q}^T\mathbf{b} + M_1(-\mathbf{n}^T\mathbf{b}, \|\mathbf{b}\|_2)\} < 1$, i.e., if $\mathbf{q}^T\mathbf{b} < 1 - M_1(\mathbf{n}^T\mathbf{b}, \|\mathbf{b}\|_2) = V_u(\mathbf{n}^T\mathbf{b}, \|\mathbf{b}\|_2)$ and $\mathbf{q}^T\mathbf{b} > -(1 - M_1(-\mathbf{n}^T\mathbf{b}, \|\mathbf{b}\|_2)) = V_l(\mathbf{n}^T\mathbf{b}, \|\mathbf{b}\|_2)$. □

APPENDIX F
PROOFS §4.4

We make use of the following lemma from [40].

Lemma 7. *If $\mathcal{R} = S(\mathbf{q}, r) \cap \{\mathbf{n}^T\boldsymbol{\theta} \leq c\}$ is nonempty, then*

$$\mathit{diam}(\mathcal{R}) = \begin{cases} 2\sqrt{r^2 - (\mathbf{n}^T\mathbf{q} - c)^2}, & \text{if } \mathbf{q} \notin \mathcal{R}; \\ 2r, & \text{otherwise.} \end{cases}$$

Lemma 4: The assumption that $0 < \psi_d \leq 1$ is equivalent to $\mathbf{q} \notin D$. Hence, under this assumption, by Lemma 7, and equations (24) and (26), the diameter of $D = D(\mathbf{q}_1, r_1; \mathbf{n}, c)$ is $2\sqrt{r_1^2 - (\mathbf{n}^T\mathbf{q}_1 - c)^2} = 2r_1\sqrt{1 - \psi_d^2} = 2r_d$. So the diameter of the circumsphere of D must be at least $2r_d$.

To show that the sphere $S(\mathbf{q}_d, r_d)$ with center $\mathbf{q}_d$ and radius r_d is the circumsphere, we show that every point $\mathbf{p}$ on the boundary of D is contained in $S(\mathbf{q}_d, r_d)$. We can write $\mathbf{p} = \mathbf{q}_d + \alpha\mathbf{v} + \beta\mathbf{n}$, where $\mathbf{v}$ is a unit norm vector in $\mathbf{n}^\perp$ and α, β are scalars with $\beta \leq 0$. We need to show that $\|\mathbf{p} - \mathbf{q}_d\|_2^2 = \alpha^2 + \beta^2 \leq r_d^2$. Since $\mathbf{p}$ is on the boundary of D, either $\beta = 0$ and $\alpha^2 \leq r_d^2$, or $\beta < 0$ and $\|\mathbf{p} - \mathbf{q}\|_2^2 = r^2$. In the first case, $\|\mathbf{p} - \mathbf{q}_d\|_2^2 = \alpha^2 \leq r_d^2$. In the second case, $r^2 = \|\mathbf{p} - \mathbf{q}\|_2^2 = \|\mathbf{q}_d - \mathbf{q} + \alpha\mathbf{v} + \beta\mathbf{n}\|_2^2 = \|(-\psi_d r + \beta)\mathbf{n} + \alpha\mathbf{v}\|_2^2 = \psi_d^2 r^2 - 2\psi_d r\beta + \alpha^2 + \beta^2$. Hence $\alpha^2 + \beta^2 = r^2(1 - \psi_d^2) + 2\beta\psi_d r < r^2(1 - \psi_d^2) = r_d^2$. □

APPENDIX G
PROOFS §4.5

Lemma 5: We first solve (16) ($m = 2$) with $\|\mathbf{b}\|_2 = 1$ by solving the Lagrangian problem:

$$\max_{\mu,\sigma,\lambda\geq 0} \min_{\boldsymbol{\theta}} \quad \mathcal{L}(\boldsymbol{\theta}, \mu, \sigma, \lambda) = -\boldsymbol{\theta}^T\mathbf{b} + \mu[(\boldsymbol{\theta} - \mathbf{q})^T(\boldsymbol{\theta} - \mathbf{q}) - r^2] + \sigma(\mathbf{n}_1^T\boldsymbol{\theta} - c_1) + \lambda(\mathbf{n}_2^T\boldsymbol{\theta} - c_2). \tag{51}$$

Solving $\partial\mathcal{L}/\partial\boldsymbol{\theta} = 0$ for $\boldsymbol{\theta}$ and substitution into $\mu_{\mathcal{R}}$ and (51) yields:

$$\begin{aligned} \mu_{\mathcal{R}}(\mathbf{b}) &= \mathbf{b}^T\mathbf{q} + \tfrac{1}{(2\mu)} - \tfrac{\sigma\mathbf{b}^T\mathbf{n}_1}{(2\mu)} - \tfrac{\lambda\mathbf{b}^T\mathbf{n}_2}{(2\mu)}, \\ \mathcal{L}(\mu, \sigma, \lambda) &= -\mathbf{q}^T\mathbf{b} - \tfrac{1}{4\mu} - \tfrac{\sigma^2}{4\mu} - \tfrac{\lambda^2}{4\mu} - \mu r^2 \\ &\quad + \sigma r\psi_1 + \lambda r\psi_2 + \tfrac{\sigma}{2\mu}t_1 + \tfrac{\lambda}{2\mu}t_2 - \tfrac{\lambda\sigma}{2\mu}\tau, \end{aligned} \tag{52}$$

where $\psi_1 = (\mathbf{q}^T\mathbf{n}_1 - c_1)/r$, $\psi_2 = (\mathbf{q}^T\mathbf{n}_2 - c_2)/r$, $t_1 = \mathbf{n}_1^T\mathbf{b}$, $t_2 = \mathbf{n}_2^T\mathbf{b}$ and $\tau = \mathbf{n}_1^T\mathbf{n}_2$. Setting the derivatives of $\mathcal{L}$ w.r.t. μ, σ and λ, respectively, to zero yields:

$$\begin{aligned} &1 + \sigma^2 + \lambda^2 - 4\mu^2r^2 = 2\sigma t_1 + 2\lambda t_2 - 2\lambda\sigma\tau \\ &\sigma = 2\mu r\psi_1 + t_1 - \lambda\tau, \quad \lambda = 2\mu r\psi_2 + t_2 - \sigma\tau. \end{aligned} \tag{53}$$

(Case I) If $\lambda = 2\mu r\psi_2 + t_2 - \sigma\tau < 0$, then set $\lambda = 0$. Substitution into (53) yields: $\sigma = 2\mu r\psi_1 + t_1$ and $1 + \sigma^2 - 4\mu^2r^2 - 2\sigma t_1 = 0$. There are two subcases:
(IA) If $t_1 > -\psi_1$, then $\sigma = t_1 + \psi_1\sqrt{\frac{1-t_1^2}{1-\psi_1^2}}$, $\mu = \frac{1}{2r}\sqrt{\frac{1-t_1^2}{1-\psi_1^2}}$ and $\lambda < 0 \Leftrightarrow (\psi_2 - \psi_1\tau)\sqrt{\frac{1-t_1^2}{1-\psi_1^2}} + t_2 - t_1\tau < 0$.
(IB) If $t_1 \leq -\psi_1$, then $\sigma = 0$, $\mu = 1/(2r)$ and $\lambda < 0 \Leftrightarrow t_2 < -\psi_2$.
(Case II) Suppose $\lambda = 2\mu r\psi_2 + t_2 - \sigma\tau > 0$. Again there are two subcases: (IIA) If $\sigma = 2\mu r\psi_1 + t_1 - \lambda\tau < 0$, then set $\sigma = 0$. Substitution into (53) yields: $\lambda = 2\mu r\psi_2 + t_2$ and $1 + \lambda^2 - 4\mu^2r^2 - 2\lambda t_2 = 0$. Solving gives,

$$\lambda = t_2 + \psi_2\sqrt{\tfrac{1-t_2^2}{1-\psi_2^2}} \quad \text{and} \quad \mu = \tfrac{1}{2r}\sqrt{\tfrac{1-t_2^2}{1-\psi_2^2}}$$

with $\lambda > 0 \Leftrightarrow t_2 > -\psi_2$ and $\sigma < 0 \Leftrightarrow (\psi_1 - \psi_2\tau)\sqrt{\frac{1-t_2^2}{1-\psi_2^2}} + t_1 - t_2\tau < 0$. (IIB) If $\sigma = 2\mu r\psi_1 + t_1 - \lambda\tau > 0$, then substituting $\lambda = 2\mu r\psi_2 + t_2 - \sigma\tau$ and $\sigma = 2\mu r\psi_1 + t_1 - \lambda\tau$ into (53) yields, $(1 - \tau^2)\sigma = 2\mu r(\psi_1 - \psi_2\tau) + t_1 - t_2\tau$ and $(1-\tau^2)\sigma^2 + 2\sigma(t_2\tau - t_1) + 4\mu^2r^2(\psi_2^2 - 1) + 1 - t_2^2 = 0$. Solving these equations gives: $\mu = \frac{1}{2r}\Delta$, $\lambda = \frac{\psi_2 - \psi_1\tau}{1-\tau^2}\Delta + \frac{t_2 - t_1\tau}{1-\tau^2} > 0$, and $\sigma = \frac{\psi_1 - \psi_2\tau}{1-\tau^2}\Delta + \frac{t_1 - t_2\tau}{1-\tau^2} > 0$, where $\Delta = \sqrt{\frac{1 + 2t_1t_2\tau - t_1^2 - t_2^2 - \tau^2}{1 + 2\psi_1\psi_2\tau - \psi_1^2 - \psi_2^2 - \tau^2}}$.

Substituting the expressions for μ, σ and λ under the various conditions into (52) yields:

[(1)] $t_1 < -\psi_1, t_2 < -\psi_2$: $\quad \mu_{\mathcal{R}}(\mathbf{b}) = \mathbf{q}^T\mathbf{b} + r$.

[(2)] $t_2 > -\psi_2, \frac{t_1 - t_2\tau}{\sqrt{1-t_2^2}} < \frac{\psi_2\tau - \psi_1}{\sqrt{1-\psi_2^2}}$:

$$\mu_{\mathcal{R}}(\mathbf{b}) = \mathbf{q}^T\mathbf{b} + r\sqrt{(1-t_2^2)(1-\psi_2^2)} - rt_2\psi_2.$$

[(3)] $t_1 > -\psi_1, \frac{t_2 - t_1\tau}{\sqrt{1-t_1^2}} < \frac{\psi_1\tau - \psi_2}{\sqrt{1-\psi_1^2}}$:

$$\mu_{\mathcal{R}}(\mathbf{b}) = \mathbf{q}^T\mathbf{b} + r\sqrt{(1-t_1^2)(1-\psi_1^2)} - rt_1\psi_1.$$

[(4)] $\frac{(t_1 - t_2\tau)}{\sqrt{1+2t_1t_2\tau - t_1^2 - t_2^2 - \tau^2}} > \frac{(\psi_2\tau - \psi_1)}{\sqrt{1+2\psi_1\psi_2\tau - \psi_1^2 - \psi_2^2 - \tau^2}}$ and $\frac{(t_2 - t_1\tau)}{\sqrt{1+2t_1t_2\tau - t_1^2 - t_2^2 - \tau^2}} > \frac{(\psi_1\tau - \psi_2)}{\sqrt{1+2\psi_1\psi_2\tau - \psi_1^2 - \psi_2^2 - \tau^2}}$:

$$\begin{aligned}\mu_{\mathcal{R}}(\mathbf{b}) = \mathbf{q}^T\mathbf{b} &- \frac{r}{1-\tau^2}((\psi_1 - \psi_2\tau)t_1 + (\psi_2 - \psi_1\tau)t_2) \\ &+ \frac{r}{1-\tau^2}\sqrt{(1-\tau^2+2\psi_1\psi_2\tau - \psi_1^2 - \psi_2^2)} \\ &\times \sqrt{(1-\tau^2+2t_1t_2\tau - t_1^2 - t_2^2)}.\end{aligned}$$

For general $\mathbf{b}$ we use $\mu_{\mathcal{R}}(\mathbf{b}) = \|\mathbf{b}\|_2\mu_{\mathcal{R}}(\mathbf{b}/\|\mathbf{b}\|_2)$. So in each of the above expressions we replace $\mathbf{b}$, $t_1 = \mathbf{n}_1^T\mathbf{b}$, and $t_2 = \mathbf{n}_2^T\mathbf{b}$ by $\mathbf{b}/\|\mathbf{b}\|_2$, $t_1/\|\mathbf{b}\|_2$ and $t_2/\|\mathbf{b}\|_2$, respectively. Then multiply each expression by $\|\mathbf{b}\|_2$. This yields the result in Lemma 5. □

Theorem 4: This is almost identical to the proof of Theorem 3 and is hence omitted. □

APPENDIX H
PROOFS §4.6

Lemma 6: Note $\max\{\mu_{\mathcal{R}}(\mathbf{b}_i), \mu_{\mathcal{R}}(-\mathbf{b}_i)\} = \max_{\boldsymbol{\theta}\in\mathcal{R}} \max\{\boldsymbol{\theta}^T\mathbf{b}, -\boldsymbol{\theta}^T\mathbf{b}\} = \max_{\boldsymbol{\theta}\in\mathcal{R}} |\boldsymbol{\theta}^T\mathbf{b}|$. If $\mathcal{R}_1$ or $\mathcal{R}_2$ is empty, the result is clear. Hence assume each is nonempty. For the lasso, if $\mathbf{b}_i$ is rejected by $T_{\mathcal{R}_1} \vee T_{\mathcal{R}_2}$, then either $\max_{\boldsymbol{\theta}\in\mathcal{R}_1} |\boldsymbol{\theta}^T\mathbf{b}_i| < 1$ or $\max_{\boldsymbol{\theta}\in\mathcal{R}_2} |\boldsymbol{\theta}^T\mathbf{b}_i| < 1$. Without loss of generality assume $\max_{\boldsymbol{\theta}\in\mathcal{R}_1} |\boldsymbol{\theta}^T\mathbf{b}_i| < 1$. Since $\mathcal{R}_1 \cap \mathcal{R}_2$ is a subset of $\mathcal{R}_1$, this implies that $\max_{\boldsymbol{\theta}\in\mathcal{R}_1\cap\mathcal{R}_2} |\boldsymbol{\theta}^T\mathbf{b}_i| \leq \max_{\boldsymbol{\theta}\in\mathcal{R}_1} |\boldsymbol{\theta}^T\mathbf{b}_i| < 1$, so $\mathbf{b}_i$ is also rejected by $T_{\mathcal{R}_1\cap\mathcal{R}_2}$. Therefore $T_{\mathcal{R}_1} \vee T_{\mathcal{R}_2} \preceq T_{\mathcal{R}_1\cap\mathcal{R}_2}$. The proof for the nn-lasso is similar. □

ACKNOWLEDGMENT

This work was partially supported by NSF grant CIF 1116208.

REFERENCES

[1] R. Tibshirani, "Regression shrinkage and selection via the lasso," *J. Royal. Statist. Soc B.*, vol. 58, no. 1, pp. 267–288, 1996.

[2] A. Y. Yang, A. Ganesh, Z. Zhou, S. S. Sastry, and Y. Ma, "A review of fast l1-minimization algorithms for robust face recognition," *Arxiv preprint arXiv:1007.3753*, 2010.

[3] M. Elad, *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*. Springer, 2010.

[4] J. Wright, Y. Ma, J. Mairal, G. Sapiro, T. Huang, and S. Yan, "Sparse representation for computer vision and pattern recognition," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 1031–1044, 2010.

[5] J. Mairal, M. Elad, and G. Sapiro, "Sparse representation for color image restoration," *IEEE Transactions on Image Processing*, vol. 17, no. 1, pp. 53–69, 2008.

[6] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman, "Non-local sparse models for image restoration," in *IEEE 12th Int. Conf. on Computer Vision*, 2009, pp. 2272–2279.

[7] J. Wright, A. Y. Yang, A. Ganesh, S. S. Sastry, and Y. Ma, "Robust face recognition via sparse representation," *IEEE Trans, on Pattern Analysis and Machine Intelligence*, vol. 31, no. 2, pp. 210–227, 2009.

[8] A. Wagner, J. Wright, A. Ganesh, Z. Zhou, H. Mobahi, and Y. Ma, "Towards a practical face recognition system: robust alignment and illumination by sparse representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2011.

[9] K. Yu, T. Zhang, and Y. Gong, "Nonlinear learning using local coordinate coding," in *Advances in Neural Information Processing Systems*, vol. 3, 2009.

[10] T. N. Sainath, A. Carmi, D. Kanevsky, and B. Ramabhadran, "Bayesian compressive sensing for phonetic classification," in *IEEE Int. Conf. Acoustics Speech and Signal Processing*, 2010, pp. 4370–4373.

[11] T. N. Sainath, B. Ramabhadran, D. Nahamoo, D. Kanevsky, and A. Sethy, "Sparse representation features for speech recognition," in *Eleventh Annual Conf. of the Int. Speech Communication Association*, 2010.

[12] K. Chang, J. Jang, and C. S. Iliopoulos, "Music genre classification via compressive sampling," in *Proc. 11th Int. Conf. on Music Information Retrieval (ISMIR)*, 2010, pp. 387–392.

[13] S. Prasad, P. Melville, A. Banerjee, and V. Sindhwani, "Emerging topic detection using dictionary learning," in *ACM Conference on Information and Knowledge Management*, 2011.

[14] D. Zhang, M. Yang, and X. Feng., "Sparse representation or collaborative representation: Which helps face recognition?" in *Int. Conf. on Computer Vision*, 2011, pp. pp. 471–478.

[15] X. Chen and P. J. Ramadge, "Collaborative representation, sparsity or nonlinearity: What is key to dictionary based classification?" in *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, May 2014.

[16] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani, "Least Angle Regression," *Annals of Statistics*, vol. 32, no. 2, pp. 407–499, 2004.

[17] J. A. Trop and A. C. Gilbert, "Signal Recovery From Random Measurements Via Orthogonal Matching Pursuit," *IEEE Transactions on Information Theory*, vol. 53, no. 12, pp. 4655–4666, Dec. 2007.

[18] S. M. Smith, "Overview of fMRI analysis," *British Journal of Radiology*, vol. 77, no. 2, p. S167, 2004.

[19] J. Fan and J. Lv, "Sure independence screening for ultrahigh dimensional feature space," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 70, no. 5, pp. 849–911, 2008.

[20] R. Tibshirani, J. Bien, J. Friedman, T. Hastie, N. Simon, J. Taylor, and R. J. Tibshirani, "Strong rules for discarding predictors in lasso-type problems," *arXiv:1011.2234*, 2010.

[21] G. Thompson, F. Tonge, and S. Zionts, "Techniques for removing nonbonding constraints and extraneous variables from linear programming problems," *Management Science*, vol. 12, no. 7, pp. 588–608, 1966.

[22] L. El Ghaoui, V. Viallon, and T. Rabbani, "Safe feature elimination in sparse supervised learning," *Pacific Journal of Optimization*, vol. 8, no. 4, pp. 667–698, 2012.

[23] Z. J. Xiang, H. Xu, and P. J. Ramadge, "Learning sparse representations of high dimensional data on large scale dictionaries," in *Advances in Neural Information Processing Systems*, 2011.

[24] Z. J. Xiang and P. J. Ramadge, "Fast lasso screening tests based on correlations," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2012.

[25] Z. J. Xiang, "Combining structural knowledge with sparsity in machine learning and signal processing," *Ph.D. Thesis, Department of Electrical Engineering*, Aug. 2012.

[26] L. Dai and K. Pelckmans, "An ellipsoidal based, two-stage screening test for bpdn," in *20th European Signal Processing Conference*, Aug 2012.

[27] J. Wang, B. Lin, P. Gong, P. Wonka, and J. Ye, "Lasso screening rules via dual polytope projection," *arXiv:1211.3966v2 [cs.LG]*, Nov. 2014.

[28] H. Zou and T. Hastie, "Regularization and variable selection via the elastic net," *Journal of the Royal Statistical Society B.*, vol. 67, no. 2, pp. 301–320, 2005.

[29] R. Tibshirani and J. Taylor, "The solution path of the generalized lasso," *Annals of Statistics*, vol. 39, no. 3, pp. 1335–1371, 2011.
[30] J.-B. Hiriart-Urruty and C. Lemarechal, *Fundamentals of Convex Analysis*. Springer, 2001.
[31] J.-J. Fuchs, "Recovery of exact sparse representations in the presence of bounded noise," *IEEE Transactions on Information Theory*, vol. 51, no. 10, pp. 3601–3608, Oct. 2005.
[32] R. Tibshirani, "The lasso problem and uniqueness," *arXiv:1206.0313 [math.ST]*, 4th Nov. 2012.
[33] P. Xu and P. J. Ramadge, "Three structural results on the lasso problem," in *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, 2013, pp. 3392–3396.
[34] L. El Ghaoui, V. Viallon, and T. Rabbani, "Safe feature elimination in sparse supervised learning," EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2010-126, Sep 2010. [Online]. Available: http://www.eecs.berkeley.edu/Pubs/TechRpts/2010/EECS-2010-126.html
[35] ——, "Safe Feature Elimination for the LASSO and Sparse Supervised Learning Problems," *arXiv:1009.4219v2 [cs.LG]*, 2011.
[36] Y. Wang, Z. Xiang, and P. Ramadge, "Tradeoffs in improved screening of lasso problems," in *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, Jun. 2013.
[37] Y. Wang, Z. J. Xiang, and P. J. Ramadge, "Lasso screening with a small regularization parameter," in *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, Jun. 2013.
[38] J. Mairal and B. Yu, "Complexity analysis of the lasso regularization path," in *Proc. of the 29th Int. Conf. on Machine Learning (ICML 2012), Edinburgh, Scotland*, 2012.
[39] S. Luo and Z. Chen, "Sequential lasso for feature selection with ultra-high dimensional feature space," *arXiv:1107.2734 [stat.ME]*, 14 July 2011.
[40] Y. Wang, X. Chen, and P. J. Ramadge, "Feedback-controlled sequential lasso screening," Princeton University, Tech. Rep., June 2013.
[41] Y. LeCun and C. Cortes, "The MNIST database of handwritten digits," 1998.
[42] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278 –2324, Nov. 1998.
[43] A. S. Georghiades, P. N. Belhumeur, and D. J. Kriegman, "From few to many: Illumination cone models for face recognition under variable lighting and pose," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 23, no. 6, pp. 643–660, 2002.
[44] K. C. Lee, J. Ho, and D. J. Kriegman, "Acquiring linear subspaces for face recognition under variable lighting," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, pp. 684–698, 2005.
[45] A. Frank and A. Asuncion, "UCI Machine Learning Repository, University of California, Irvine, School of Information and Computer Sciences," 2010. [Online]. Available: http://archive.ics.uci.edu/ml
[46] H. Lee, A. Battle, R. Raina, and A. Ng, "Efficient sparse coding algorithms," in *Advances in Neural Information Processing Systems*, vol. 19, 2007, p. 801.
[47] P. S. and J. Theiler, "Online feature selection using grafting," in *Proc. of the Int. Conf. on Machine Learning*, 2003, pp. 592–599.
[48] J. Liu, Z. Zhao, J. Wang, and J. Ye, "Safe screening with variational inequalities and its application to lasso," *arXiv:1307.7577v2 [cs.LG]*, Oct. 2013.
[49] S. Dash, D. Malioutov, and K. R. Varshney, "Screening for learning classification rules via boolean compressed sensing," in *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, May. 2014.
[50] Y. Wang, X. Chen, and P. J. Ramadge, "Sparse representation classification via sequential lasso screening," in *1st IEEE Global Conference on Signal and Information Processing*, Dec. 2013.
[51] P.-K. Jao, C. C. M. Yeh, and Y.-H. Yang, "Modified lasso screening for audio word-based music classification using large scale dictionary," in *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, May. 2014.
[52] X. Chen and P. J. Ramadge, "Music genre classification using multiscale scattering and sparse representations," in *Proc. Annual Conference on Information Sciences and Systems, Johns Hopkins Univ.*, 2013, pp. 1–6.